

# Robust causal inference with continuous instruments using the local instrumental variable curve

Edward H. Kennedy \*

Department of Statistics, Carnegie Mellon University

Scott A. Lorch

Department of Pediatrics, Perelman School of Medicine,  
University of Pennsylvania

Dylan S. Small

Department of Statistics, The Wharton School, University of Pennsylvania

## Abstract

Instrumental variables are commonly used to estimate effects of a treatment afflicted by unmeasured confounding, and in practice instruments are often continuous (e.g., measures of distance, or treatment preference). However, available methods for continuous instruments have important limitations: they either require restrictive parametric assumptions for identification, or else rely on modeling both the outcome and treatment process well (and require modeling effect modification by all adjustment covariates). In this work we develop the first semiparametric doubly robust estimators of the local instrumental variable effect curve, i.e., the effect among those who would take treatment for instrument values above some threshold and not below. In addition to being robust to misspecification of either the instrument or treatment/outcome processes, our approach also incorporates information about the instrument mechanism and allows for flexible data-adaptive estimation of effect modification. We discuss asymptotic properties under weak conditions, and use the methods to study infant mortality effects of neonatal intensive care units with high versus low technical capacity, using travel time as an instrument.

*Keywords:* causal inference, complier average treatment effect, cross-validation, doubly robust, semiparametric theory.

---

\*Edward Kennedy is Assistant Professor of Statistics, Carnegie Mellon University, Pittsburgh, PA 15217 (e-mail: edward@stat.cmu.edu). Scott Lorch is Associate Professor of Pediatrics at the University of Pennsylvania School of Medicine and Attending Neonatologist at The Children's Hospital of Philadelphia. Dylan Small is Professor of Statistics, University of Pennsylvania. The authors gratefully acknowledge support from NIH grant R01-DK090385 (Kennedy) and NSF grant SES-1260782 (Small).

# 1 Introduction

Instrumental variables (IVs) are commonly used to estimate effects of treatments that are afflicted by unmeasured confounding. Instruments are special variables that influence treatment, but are themselves unconfounded and do not directly affect outcomes, allowing the recovery of some causal information from data that might otherwise be unusable. In practice, instruments are often continuous (e.g., measures of distance, or treatment preference), but most available methods only consider instruments that are discrete (and typically binary). Further, methods that do allow for continuous IVs have important limitations.

Classical IV methods (e.g., standard two-stage least squares), which were developed in a structural equation model framework, allow for continuous instruments but require strong parametric assumptions for identification, assume that treatment effects do not vary across units, and also require correct parametric models for at least how the outcome process depends on covariates and instruments (Wooldridge 2010; Okui *et al.* 2012). Alternatively, Robins and others (Robins 1989; Robins 1994; Hernán & Robins 2006; Tan 2010; Okui *et al.* 2012) developed approaches in the potential outcomes framework that can also handle continuous instruments, but which allow heterogeneous treatment effects, and also permit doubly robust covariate adjustment (Robins & Rotnitzky 2001; van der Laan & Robins 2003; Bang & Robins 2005). Doubly robust IV methods are consistent as long as either the instrument mechanism or the treatment/outcome mechanisms are correctly modeled (not necessarily both), and they can also yield fast root-n convergence rates and inference even when using flexible nonparametric methods for covariate adjustment. However, the methods developed in this framework still require parametric assumptions for identification; they typically target treatment effects on the treated, and achieve identification with dimension-reducing parametric assumptions that restrict how heterogeneous treatment effects can be. As noted for example by Tchetgen Tchetgen & Vansteelandt (2013), this kind of approach is problematic because a priori information about the parametric form of underlying causal structure is rarely available, and misspecification could lead to large biases that cannot be detected with data.

An alternative approach is to replace dimension-reducing homogeneity assumptions with a monotonicity assumption (Robins 1989; Imbens & Angrist 1994), which rules out the possibility that any units would respond oppositely to encouragement from the instrument. In other words, there can be units who are encouraged by the instrument, as well as units who do not respond at all to the instrument, but there cannot be units who defy encouragement from the instrument. For example, in the binary instrument case, there can be units who take treatment if and only if they receive the instrument, as well as units who always or never take treatment; however, there cannot be units who take control if the instrument is received but take treatment if not. This assumption is often plausible in practice, and also permits nonparametric identification of causal effects among compliers (i.e., those who do respond to encouragement from the instrument). However, monotonicity is usually framed in terms of binary instruments (Imbens & Angrist 1994; Abadie 2003; Tan 2006; Ogburn *et al.* 2015). An important exception is a strand of work that has focused on estimating local IV (LIV) curves, i.e., effects among units who would comply right at a given threshold value of the instrument (Heckman 1997; Heckman & Vytlacil

1999; Glickman & Normand 2000; Heckman & Vytlacil 2005).

This literature on LIV approaches arose out of a latent index or selection model framework (Vytlacil 2002), and is unique in allowing for continuous instruments while still permitting nonparametric identification. However there are important limitations. First, available approaches for estimating the LIV curve rely on modeling how both the treatment and outcome depend on covariates and instrument (Basu *et al.* 2007; Carneiro *et al.* 2010), and typically use restrictive parametric models. This is problematic since parametric models are often relied upon based on convenience, rather than real substantive knowledge, and can yield severe bias if misspecified. Conversely, fully nonparametric approaches are sensitive to the curse of dimensionality and typically yield estimators with slow rates of convergence and no hope for confidence intervals (without impractical undersmoothing). Further, in the IV setting there may be some information available about how the instrument depends on covariates (e.g., about the density of the instrument given covariates), but this is not incorporated in approaches that rely solely on treatment and outcome models, whether parametric or nonparametric. Our paper solves these problems with a semiparametric doubly robust approach that can attain parametric rates of convergence, even while allowing flexible nonparametric estimation of nuisance functions.


Second, available LIV estimands are fully conditional on all measured covariates, even though in many cases effect modification is not of particular scientific interest, or else it is only of interest for a small subset of covariates. Marginal effects are less ambitious parameters that can be estimated more robustly and are often more closely tied to scientific questions. van der Laan & Robins (2003) point out that using fully conditional effects puts us at the whim of whatever confounders happen to arise in the dataset at hand, whereas marginal effects allow framing scientific questions a priori. Marginalizing currently available fully conditional estimators leads to awkward and uninterpretable models, as discussed in the local average treatment effect (LATE) setting by Ogburn *et al.* (2015). However, in contrast to the LATE setting, currently available methods cannot be adapted for doubly robust estimation of the LIV curve, even in the simpler case where effect modifiers are discrete or not present. Our paper solves these problems by developing methods for working models of a marginal version of the LIV curve itself, allowing for arbitrary effect modification. Our direct modeling approach eases interpretability by allowing analysts to incorporate background knowledge on the actual parameter of interest, without having to specify models for nuisance quantities that are not of direct scientific interest.

In addition to the above, our work makes several other important advances. Importantly, we use empirical process theory and sample splitting to derive asymptotic properties of our approach under weak conditions, which allow for flexible data-adaptive estimation of nuisance functions in the presence of complex high-dimensional confounding. We also develop a doubly robust cross-validation approach for model selection in high-dimensional settings, which is crucial for learning the LIV curve from data. Finally, we explore finite-sample properties via simulation, and implement our methods to study effects of high-level neonatal intensive care units (NICUs) on infant mortality, using travel time as an instrument.

## 2 Preliminaries

### 2.1 Data & Notation

Suppose we observe an independent and identically distributed sample  $(\mathbf{O}_1, \dots, \mathbf{O}_n)$  with  $\mathbf{O} = (\mathbf{X}, Z, A, Y)$ , where  $\mathbf{X}$  is a vector of covariates,  $Z$  is a continuous instrument for a binary treatment  $A$ , and  $Y$  is some real-valued outcome of interest. The covariates  $\mathbf{X} = (\mathbf{V}, \mathbf{W})$  are partitioned into potential effect modifiers of interest  $\mathbf{V}$  and other covariates  $\mathbf{W} = \mathbf{X} \setminus \mathbf{V}$  not of interest but for which adjustment is still necessary. The choice of  $\mathbf{V}$  is based purely on the scientific question, so that if effect modification is not of interest one can simply select  $\mathbf{V} = \emptyset$ . We characterize causal effects using potential outcome notation (Rubin 1974), and so let  $Y^a$  (and  $Y^{za}$ ) denote the potential outcomes that would have been observed had treatment level  $A = a$  (and instrument level  $Z = z$ ) been received. Similarly we let  $A^z$  denote the potential treatment that would have been observed under instrument level  $Z = z$ . A directed acyclic graph showing the data structure is given in Figure 1.



**Figure 1:** Directed acyclic graph showing covariates  $\mathbf{X}$  (consisting of potential effect modifiers of interest  $\mathbf{V}$  and other variables  $\mathbf{W}$ ), instrument  $Z$ , treatment  $A$ , outcome  $Y$ , and unmeasured variables  $U$ . Gray dotted arrows indicate relationships that will be assumed absent by identifying assumptions.

We let  $P$  denote the distribution of  $\mathbf{O}$ , with density with respect to some dominating measure given by  $p(\mathbf{o}) = p(y | \mathbf{x}, z, a)p(a | \mathbf{x}, z)p(z | \mathbf{x})p(\mathbf{x})$ . In general we write the density of a variable  $T$  under  $P$  evaluated at  $z$  as  $p(T = z)$ , except when there is no ambiguity (e.g.,  $p(t)$  is the density of  $T$  at  $t$ ), and we use  $\text{supp}(T)$  for the support of  $T$ . Finally we use some additional notation to simplify the presentation. Specifically we let  $\pi(z | \mathbf{x}) = p(Z = z | \mathbf{X} = \mathbf{x})$  denote the density of the instrument given covariates (i.e., instrument propensity score),  $\mu(\mathbf{x}, z) = \mathbb{E}(Y | \mathbf{X} = \mathbf{x}, Z = z)$  and  $\lambda(\mathbf{x}, z) = \mathbb{E}(A | \mathbf{X} = \mathbf{x}, Z = z)$  denote the outcome and treatment regression functions, respectively, with marginalized versions given by  $m(z, \mathbf{v}) = \mathbb{E}\{\mu(\mathbf{X}, z) | \mathbf{V} = \mathbf{v}\}$  and  $\ell(z, \mathbf{v}) = \mathbb{E}\{\lambda(\mathbf{X}, z) | \mathbf{V} = \mathbf{v}\}$ . We let  $\mathbb{P}_n$  denote the empirical measure so that empirical averages can be written as  $n^{-1} \sum_i f(\mathbf{O}_i) = \mathbb{P}_n\{f(\mathbf{O})\}$ . The notation  $\|\cdot\|$  denotes the Euclidean norm  $\|\boldsymbol{\beta}\| = (\boldsymbol{\beta}^T \boldsymbol{\beta})^{1/2}$ , and  $\|f\|_2 = \{\int f(\mathbf{o})^2 dP(\mathbf{o})\}^{1/2}$  denotes the  $L_2(P)$  norm.

## 2.2 Monotonicity

Before defining the causal estimand of interest and considering identifying assumptions, it will be helpful to discuss the concept of monotonicity, which was introduced in various forms by Robins (1989) and Imbens & Angrist (1994), among others. In a classical binary instrument setting (where  $\text{supp}(Z) = \{0, 1\}$ ), monotonicity can be stated succinctly as

$$A^1 \geq A^0 \text{ with probability one.}$$

Monotonicity rules out the possibility that there are troublesome units in the population with  $A^0 = 1$  but  $A^1 = 0$ . Such units are called ‘defiers’ since they take treatment  $A = 1$  when not encouraged by the instrument (i.e., when  $Z = 0$ ), but take control  $A = 0$  when they are in fact encouraged (i.e., when  $Z = 1$ ). Thus monotonicity ensures that the population only comprises never-takers ( $A^0 = A^1 = 0$ ), always-takers ( $A^0 = A^1 = 1$ ), and compliers ( $A^0 = 0, A^1 = 1$ ). Monotonicity can often be a reasonable assumption in practice, but not always; it has been discussed extensively in previous work, particularly for binary instruments (see Imbens (2014) and discussion for a nice overview).

A natural way to extend monotonicity to the continuous instrument setting is as follows.

*Assumption 1* (Monotonicity). If  $z' > z$  then  $A^{z'} \geq A^z$  with probability one.

Under Assumption 1, no unit would ever change from treatment to control with an increase in the instrument value; increasing the instrument can either encourage treatment over control or have no effect at all, but it cannot discourage treatment relative to lesser instrument values. Thus the population still comprises never-takers, always-takers, and compliers, but with continuous instruments the compliers can be further partitioned into compliers at given instrument values. In particular, a complier at  $Z = z$  would be a unit for which  $A^z = 1$  but  $A^{z-\delta} = 0$  for any  $\delta > 0$ . We use a population version of monotonicity for simplicity, but only a covariate-specific conditional version is necessary when targeting conditional treatment effects (Tan 2006). The above continuous version of monotonicity has been employed and discussed by Glickman & Normand (2000) and Vytlačil (2002), for example. Importantly, these authors showed that (when coupled with standard identifying assumptions to be discussed shortly) the above monotonicity assumption is equivalent to the following latent threshold model.

*Assumption 1'* (Latent Threshold).  $A^z = \mathbb{1}(z \geq T)$  for an unobserved random threshold  $T$ .

Under the latent threshold model, each complier has some instrument value at which they are encouraged to take treatment, while for any lesser value they would take control. Larger values of the threshold  $T$  indicate units that are less willing to take treatment, i.e., less susceptible to encouragement by the instrument. We can thus define the latent threshold  $T$  as

$$T = \begin{cases} -\infty & \text{if } A^z = 1 \text{ for all } z \text{ (always-takers)} \\ \inf\{z : A^z = 1\} & \text{if } A^{z'} > A^z \text{ for some } z' > z \text{ (compliers)} \\ \infty & \text{if } A^z = 0 \text{ for all } z \text{ (never-takers).} \end{cases}$$

We refer to Vytlačil (2002) for further discussion and detail.

## 2.3 Estimand & Identification

In this paper our goal is estimation and inference for the local instrumental variable (LIV) curve, which we define as

$$\gamma(t, \mathbf{v}) = \mathbb{E}(Y^1 - Y^0 \mid T = t, \mathbf{V} = \mathbf{v}). \quad (1)$$

This is the average treatment effect among those with latent threshold  $T = t$  (and value  $\mathbf{V} = \mathbf{v}$  for some arbitrary baseline covariate subset), i.e., the effect among those units with  $\mathbf{V} = \mathbf{v}$  who would be encouraged to take treatment right when the instrument passes  $Z = t$  but not for lesser values. A fully conditional version of the LIV curve with  $\mathbf{V} = \mathbf{X}$  was proposed in the latent index or selection model framework by Heckman (1997), and discussed in detail by Heckman & Vytlacil (1999) and Heckman & Vytlacil (2005). Note that the variables  $\mathbf{V}$  are arbitrary, and no causal assumptions need to be made about them; this is because we are estimating effect modification (how effects vary with  $\mathbf{V}$ ) and not mediation (how effects would change if we set  $\mathbf{V}$  to certain values).

We study the LIV curve for continuous IVs since such IVs are very common in practice and since novel theory is required; an analogous LIV effect  $\mathbb{E}(Y^1 - Y^0 \mid A^t > A^{t-1}, \mathbf{V})$  for discrete multivalued instruments could be studied by adapting results from Tan (2006) and Ogburn *et al.* (2015). However, for continuous instruments, the LIV effect curve is substantially different from the local average treatment effect (LATE) studied by Angrist *et al.* (1996), Tan (2006), Ogburn *et al.* (2015) and others. Specifically, the continuous IV extension of the LATE is the three-dimensional effect surface  $\mathbb{E}(Y^1 - Y^0 \mid A^z > A^{z'}, \mathbf{V} = \mathbf{v})$  for values  $z > z'$ , which represents the effect among those who would take treatment at  $Z = z$  but not  $Z = z'$ . In contrast, for continuous instruments, the LIV parameter in (1) is an arguably more interpretable (and easier to model) curve representing the effect among those “at the margin” who would take treatment at  $Z = t$  but not  $Z < t$ . These and other differences between the LATE and LIV are discussed in more detail by Heckman & Vytlacil (1999) and Heckman & Vytlacil (2005). Further, even more substantial differences arise after identification, as discussed below.

*Remark 1.* In Heckman’s selection model framework, the LIV curve in (1) with  $\mathbf{V} = \mathbf{X}$  was termed the “marginal treatment effect”, and its observed data counterpart the “local IVs” estimand (after employing identifying assumptions). We use the latter in both cases, since “marginal” is often used to mean “averaged” instead of “at the margin”.

For identification we rely on standard IV assumptions, which have been employed for example by Angrist *et al.* (1996), Tan (2006), Ogburn *et al.* (2015), and others; useful overviews and discussions are given by Hernán & Robins (2006), Imbens (2014) (with discussion), and Baiocchi *et al.* (2014).

*Assumption 2* (Consistency).  $A = A^Z$  and  $Y = Y^A$  with probability one.

*Assumption 3* (Positivity).  $(z, \mathbf{x}) \in \text{supp}(Z, \mathbf{X})$  if  $\mathbf{x} \in \text{supp}(\mathbf{X})$ .

*Assumption 4* (Unconfoundedness of  $Z$ ).  $(Y^z, A^z) \perp\!\!\!\perp Z \mid \mathbf{X}$ .

*Assumption 5* (Exclusion Restriction).  $Y^{za} = Y^a$  with probability one.

Consistency means potential treatments  $A^z$  and outcomes  $Y^a$  are uniquely defined by a unit's own instrument and treatment levels, respectively, and not by others' levels (i.e., no interference), and also not by the way the instrument or treatment are administered (i.e., no different versions). Positivity says that the instrument is not deterministic, in the sense that every unit has some chance of receiving each level of the instrument, regardless of covariates. Unconfoundedness says that the instrument is essentially randomized once we condition on covariates, i.e., that it is unrelated to potential outcomes  $Y^z = Y^{zA^z}$  and treatments  $A^z$  under different instrument values  $Z = z$ . The exclusion restriction says that the instrument only affects outcomes through treatment. Assumptions 2–5 can hold by design in trials where the instrument is externally randomized by investigators, but in observational studies these assumptions are typically untestable and require justification based on subject matter.

Finally we also employ the following regularity conditions on the latent threshold distribution and LIV curve.

*Assumption 6* (Instrumentation).  $\inf_t p(t \mid \mathbf{v}) > 0$ .

*Assumption 7* (Continuity).  $T$  is continuously distributed and  $\gamma(t, \mathbf{v})$  is continuous in  $t$ .

Instrumentation means there are at least some units whose treatment would be affected by the value of the instrument (i.e., some who would take treatment when the instrument passes  $Z = t$ ); for now we leave the set over which the infimum is taken ambiguous.

*Remark 2.* A caveat is in order regarding Assumptions 3 and 6: positivity and instrumentation are particularly strong requirements for continuous instruments. Positivity may be violated, for example, if some units simply have no chance at receiving IV levels far away from what they actually received; e.g., in the analysis in Section 5, it may be implausible that some patients who actually live very close to high-level ICUs could ever live very far from them. Similarly, instrumentation can be violated with weak instruments, since moving the IV past certain levels simply may not affect any units' treatment status. Positivity violations can be partially ameliorated by an appropriate choice of weight function (as discussed in Section 3.1), and instrumentation violations by an appropriate choice of the set  $\mathcal{T}$  on which the LIV curve is to be estimated; however neither is a panacea. See Petersen *et al.* (2010), Westreich & Cole (2010), Stock *et al.* (2012), and Baiocchi *et al.* (2014) and others for relevant discussion of these issues in related settings. We will pursue extensions dealing with violations of these assumptions in future work.

The following theorem indicates that the LIV curve can be identified with observed data, under the above assumptions.

**Theorem 1.** *Suppose Assumption 1' holds. Let  $\mathcal{T} \subset \text{supp}(Z)$  denote a compact set on which we wish to identify  $\gamma(t, \mathbf{v})$ . If Assumptions 2–5 hold for all  $z \in \mathcal{T}$  and Assumptions 6–7 hold for all  $t \in \mathcal{T}$ , then the LIV curve is identified for any  $t \in \mathcal{T}$  by*

$$\gamma(t, \mathbf{v}) = \frac{\frac{\partial}{\partial z} \mathbb{E}\{\mathbb{E}(Y \mid \mathbf{X}, Z = z) \mid \mathbf{V} = \mathbf{v}\}}{\frac{\partial}{\partial z} \mathbb{E}\{\mathbb{E}(A \mid \mathbf{X}, Z = z) \mid \mathbf{V} = \mathbf{v}\}} \Big|_{z=t}. \quad (2)$$

A proof of Theorem 1 is given in the Supplementary Materials; the logic follows as in more standard settings where  $Z$  is discrete and  $\mathbf{V} = \mathbf{X}$ . Importantly, the LIV curve can only be identified on subsets of  $\text{supp}(Z)$ ; thus as in the binary instrument setting, we cannot identify effects for never-takers or always-takers with  $T = \pm\infty$ . As discussed in more detail in the next section, the ratio-of-derivatives structure of the LIV curve makes its study particularly interesting from a theoretical perspective, especially relative to the LATE parameter studied by Ogburn *et al.* (2015) and others.

*Remark 3.* 1 From this point forward, when we write  $\gamma(t, \mathbf{v})$  we mean the observed data expression in (2), which represents the causal effect given in (1) under Assumptions 2–7 as described in Theorem 1. Of course, if the conditions of Theorem 1 do not hold then the observed data expression in (2) may represent something other than the aforementioned causal effect. For example, if only Assumptions 2–4 hold, then we can only think of the instrument as an unconfounded continuous exposure (or dose), and  $\gamma(t, \mathbf{v})$  would represent the ratio of derivatives of the dose-response curves  $\mathbb{E}(Y^z \mid \mathbf{V} = \mathbf{v})$  and  $\mathbb{E}(A^z \mid \mathbf{V} = \mathbf{v})$ .

### 3 Main Results

In this section we develop semiparametric theory for models of the LIV curve defined in (2), use this theory to develop novel estimators (including inverse-probability-weighted, regression, and doubly robust estimators), describe asymptotic properties, and finally present cross-validation methods for model selection in high-dimensional settings.

#### 3.1 Semiparametric Theory

Suppose we have a parametric model for the LIV curve, which we write as  $\gamma(t, \mathbf{v}; \boldsymbol{\psi})$  for some finite-dimensional  $\boldsymbol{\psi} \in \mathbb{R}^q$ . Importantly, we do not assume this model is necessarily correct, and instead follow Neugebauer & van der Laan (2007), Rosenblum & van der Laan (2010), and others in using a working model approach, by formulating our estimand as the projection of the true curve  $\gamma(t, \mathbf{v})$  onto the posed working model. Specifically, we use the weighted least squares projection given by

$$\boldsymbol{\psi}_0 = \arg \min_{\boldsymbol{\psi} \in \mathbb{R}^q} \mathbb{E} \left[ w(T, \mathbf{V}) \{ \gamma(T, \mathbf{V}) - \gamma(T, \mathbf{V}; \boldsymbol{\psi}) \}^2 \right], \quad (3)$$

where  $w(t, \mathbf{v})$  is some user-specified weight function. We use the  $L_2$  loss-based projection for its convenience and familiarity, but it would be worthwhile to develop results for other loss functions in future work. Whether to use  $\gamma(t, \mathbf{v}; \boldsymbol{\psi})$  only for projections or to assume it is true is a bias-variance trade-off. If the model happens to be correct, then the projection approach will result in appropriate but generally not fully efficient estimators. However, if the posited model is incorrect, then the projection is still validly defined as a best-fitting wrong model, while a model-based approach would technically no longer be applicable and can be more difficult to interpret. The projection approach essentially formalizes how models are often viewed as approximations in practice.

An important issue in using projections is choosing the weight  $w(t, \mathbf{v})$ . Ideally there will be some subject matter justification for learning about specific parts of the effect curve. If



not, in theory one could use a uniform weight that assigns mass equally across the support, but in practice this could lead to poor efficiency. We have found that weights based on the instrument density, e.g.,  $w(t, \mathbf{v}) = p(z = t, \mathbf{v})$  or  $w(t, \mathbf{v}) = w(t) = p(z = t)$ , work well when no particular weight function is preferred based on substantive concerns. Of course if  $\gamma$  is correctly specified, all weights yield consistent estimators that only vary in efficiency.

Note that the projection parameter in (3) depends on the distribution of the latent threshold  $T$ . Although this threshold is not observed directly, its distribution is identified in the observed data (under Assumptions 1–7). For example when  $t \in \text{supp}(Z)$  we have

$$p(t \mid \mathbf{v}) = \frac{\partial}{\partial z} \mathbb{E}\{\mathbb{E}(A \mid \mathbf{X}, Z = z) \mid \mathbf{V} = \mathbf{v}\} \Big|_{z=t} \quad (4)$$

Importantly, the above expression for the threshold density equals the denominator of  $\gamma(t, \mathbf{v})$  given in Theorem 1.

After characterizing the parameter of interest in terms of observed data as in (3), based on the expression in (2), it is possible to estimate it using any number of approaches, such as parametric or nonparametric maximum likelihood, or Bayesian methods. In our setting, however, semiparametric approaches have a number of important advantages. First, they can incorporate information about the instrument mechanism, which may be better understood or easier to model than the outcome and treatment mechanisms (which is what a likelihood-based approach would rely on modeling). Second, they allow for double robustness, which means consistent estimation of  $\psi$  is possible as long as either the instrument mechanism or the treatment/outcome mechanisms are correctly modeled (not necessarily all three, so either the instrument or the treatment and outcome models can be misspecified). And third, semiparametric doubly robust approaches allow for fast root- $n$  rates of convergence for the parameter of interest  $\psi$ , even when nuisance functions are estimated at slower rates, e.g., using flexible data-adaptive or machine learning methods. Thus these estimators are less sensitive to the curse of dimensionality; this phenomenon was noted recently for example by van der Vaart (2014) and Chernozhukov *et al.* (2016), who refer to it as orthogonality. Nonetheless our results also lead to novel inverse-weighting and regression-based estimators; the latter in particular might be preferable in small samples.

A crucial aspect of developing semiparametric theory and corresponding estimators for a given problem involves characterizing the possible influence functions, and in particular finding the efficient influence function. Many details on semiparametric theory are available elsewhere (Bickel *et al.* 1993; van der Laan & Robins 2003; Tsiatis 2006; Kennedy 2016), so we give only a brief review here. Any regular asymptotically linear estimator minus its target parameter can be expressed as the empirical average of its so-called influence function plus an  $o_p(1/\sqrt{n})$  error term. Viewed as elements of a Hilbert space of mean-zero finite-variance functions equipped with covariance norm, the influence functions under a given model lie in the orthogonal complement of the nuisance tangent space. The efficient influence function can then be defined as the influence function with smallest variance, the projection of any influence function onto the tangent space of scores, or as a particular pathwise derivative. The efficient influence function is especially important in practice because its variance is the semiparametric efficiency bound (thus providing a benchmark for efficient estimation), and because it can be used to construct estimators that are doubly robust and potentially semiparametric efficient.

A major challenge in deriving semiparametric theory for the projection parameter in (3) is its complexity; namely, it is a weighted projection of a ratio of derivatives of regression functions that are partially marginalized. We conjecture that this may be why a doubly robust estimator for the LIV curve has yet to appear in the literature. In general, such complex structure would yield a complicated efficient influence function involving derivatives of regression functions, making corresponding estimators very difficult to compute. A major contribution of our work is an expression for the efficient influence function that only involves derivatives based on the known (and generally more simple) model and weight functions, rather than unknown complex and high-dimensional nuisance functions. Although still complex, our formulation yields efficient estimators that are easier to construct in practice. The next theorem gives the efficient influence function for the parameters of the LIV curve projection.

**Theorem 2.** *Suppose the weight function  $w(t, \mathbf{v})$  is continuously differentiable in  $t$  and satisfies  $w(t, \mathbf{v}) = 0$  for  $t \notin \text{int}(\mathcal{T})$ , with the set  $\mathcal{T} \subset \text{supp}(Z)$  defined as in Theorem 1. Also assume that partial derivatives (with respect to  $\boldsymbol{\psi}$  and  $t$ ) of the working model  $\gamma(t, \mathbf{v}; \boldsymbol{\psi})$  exist and are continuous. Then, under a nonparametric model, the efficient influence function for  $\boldsymbol{\psi}$  defined in (2) and (3) is proportional to*

$$\begin{aligned} \boldsymbol{\varphi}(\mathbf{O}; \boldsymbol{\psi}, \boldsymbol{\eta}) = & \mathbf{g}_1(Z, \mathbf{V}; \boldsymbol{\psi}) \left\{ \frac{A - \mathbb{E}(A \mid \mathbf{X}, Z)}{p(Z \mid \mathbf{X})} \right\} - \mathbf{g}_2(Z, \mathbf{V}; \boldsymbol{\psi}) \left\{ \frac{Y - \mathbb{E}(Y \mid \mathbf{X}, Z)}{p(Z \mid \mathbf{X})} \right\} \quad (5) \\ & + \int_{\mathcal{T}} \left\{ \mathbf{g}_1(t, \mathbf{V}; \boldsymbol{\psi}) \mathbb{E}(A \mid \mathbf{X}, Z = t) - \mathbf{g}_2(t, \mathbf{V}; \boldsymbol{\psi}) \mathbb{E}(Y \mid \mathbf{X}, Z = t) \right\} dt, \end{aligned}$$

where  $\boldsymbol{\eta} = (\pi, \lambda, \mu)$  denotes the nuisance functions defined in Section 2.1, and  $\mathbf{g}_1$  and  $\mathbf{g}_2$  are the  $(q \times 1)$  vectors

$$\begin{aligned} \mathbf{g}_1(z, \mathbf{v}; \boldsymbol{\psi}) &= \frac{\partial}{\partial t} \left\{ \frac{\partial}{\partial \boldsymbol{\psi}^*} \gamma(t, \mathbf{v}; \boldsymbol{\psi}^*) \Big|_{\boldsymbol{\psi}^* = \boldsymbol{\psi}} w(t, \mathbf{v}) \gamma(t, \mathbf{v}; \boldsymbol{\psi}) \right\} \Big|_{t=z} \\ \mathbf{g}_2(z, \mathbf{v}; \boldsymbol{\psi}) &= \frac{\partial}{\partial t} \left\{ \frac{\partial}{\partial \boldsymbol{\psi}^*} \gamma(t, \mathbf{v}; \boldsymbol{\psi}^*) \Big|_{\boldsymbol{\psi}^* = \boldsymbol{\psi}} w(t, \mathbf{v}) \right\} \Big|_{t=z}. \end{aligned}$$

A proof of Theorem 2 is given in the Supplementary Materials, and proceeds by showing that the function  $\boldsymbol{\varphi}$  is the canonical gradient of the pathwise derivative of  $\boldsymbol{\psi}$ . Importantly, our derivation of the efficient influence function uses integration by parts to transfer derivatives of the partially marginalized treatment/outcome regression functions to the known model  $\gamma(t, \mathbf{v}; \boldsymbol{\psi})$  and weight function  $w(t, \mathbf{v})$ . As noted above, this means the efficient influence function can be evaluated without analytical differentiation of the regression functions, which makes it much more practical for constructing and implementing estimators, as discussed in more detail in the next subsection. The user-specified weight function is required to vanish outside the interior of the set  $\mathcal{T} \subset \text{supp}(Z)$  since the LIV curve  $\gamma(t, \mathbf{v})$  is not identified outside of  $\mathcal{T}$  as discussed in Theorem 1. Roughly speaking, the efficient influence function  $\boldsymbol{\varphi}$  can be viewed as consisting of inverse-probability-weighted terms (the added term in the first line in (5)) plus an augmentation term (the second line and subtracted terms in the first line in (5)). This follows the general structure of influence functions in more common causal inference and missing data problems, although the form of the functions  $\mathbf{g}_1$  and  $\mathbf{g}_2$  makes the expression non-standard.

### 3.2 Proposed Method

Once we have derived the efficient influence function, we can use it to construct estimators that have numerous advantageous properties. A standard approach is to solve an estimating equation based on an estimated version of the efficient influence function; specifically we can use  $\varphi$  as an estimating function, with unknown nuisance functions replaced with estimates.

Thus our proposed estimator for a given working model  $\gamma(t, \mathbf{v}; \boldsymbol{\psi})$  is given by  $\hat{\boldsymbol{\psi}}$ , defined as the solution in  $\boldsymbol{\psi}$  to the estimating equation

$$\mathbb{P}_n\{\varphi(\mathbf{O}; \boldsymbol{\psi}, \hat{\boldsymbol{\eta}})\} = \mathbf{0}, \quad (6)$$

where  $\hat{\boldsymbol{\eta}} = (\hat{\pi}, \hat{\lambda}, \hat{\mu})$  are estimated versions of the three nuisance functions  $\boldsymbol{\eta} = (\pi, \lambda, \mu)$ . An inverse-probability-weighted estimator can be obtained by using  $\hat{\lambda} = \hat{\mu} = 0$ , and a regression estimator can be obtained by using  $\hat{\pi} = \infty$ . Another option for constructing estimators based on influence functions is targeted minimum loss-based methodology (van der Laan & Rubin 2006), which yields plug-in estimators that respect the bounds of the parameter space. In our setting, our proposed estimating equation approach will also respect any such bounds, as long as the chosen working model does; it is also relatively straightforward to implement.

*Remark 4.* For notational simplicity, the estimator proposed above uses nuisance estimates constructed from the entire sample (with asymptotic results in Section 3.3 relying on empirical process conditions). In Section 4 of the Supplementary Materials we present a sample-splitting estimator in the same spirit as Zheng & van der Laan (2010) and Chernozhukov *et al.* (2016), which does not require empirical process conditions and accommodates arbitrarily complex and adaptive nuisance estimators, e.g., random forests.

First consider the simple case where  $\mathbf{V} = \emptyset$  (i.e., effect modification is not of interest), and the LIV curve  $\gamma(t)$  is projected onto a constant  $\gamma(t; \boldsymbol{\psi}) = \psi$ . Here the target estimand  $\psi$  is a simple weighted average of  $\gamma(t)$ , of the form

$$\psi = \int_{\mathcal{T}} w^*(t) \gamma(t) dt$$

with weight  $w^*(t) = w(t)p(t) / \int_{\mathcal{T}} w(t)p(t) dt$ , where  $w(t) = w(t, \mathbf{v})$  since  $\mathbf{V}$  is empty. The quantity  $\psi$  can also be viewed as the mean treatment effect (among compliers) in a population where the density of the latent threshold  $T$  among compliers equals  $w^*(t)$ . Then solving (6) leads to the ratio estimator

$$\hat{\psi} = \frac{\int_{\mathcal{T}} w'(t) \hat{m}(t) dt + \mathbb{P}_n \left\{ w'(Z) \frac{Y - \hat{\mu}(\mathbf{X}, Z)}{\hat{\pi}(Z|\mathbf{X})} \right\}}{\int_{\mathcal{T}} w'(t) \hat{\ell}(t) dt + \mathbb{P}_n \left\{ w'(Z) \frac{A - \hat{\lambda}(\mathbf{X}, Z)}{\hat{\pi}(Z|\mathbf{X})} \right\}} \quad (7)$$

where  $\hat{m}(t) = \mathbb{P}_n\{\hat{\mu}(\mathbf{X}, t)\}$  and  $\hat{\ell}(t) = \mathbb{P}_n\{\hat{\lambda}(\mathbf{X}, t)\}$  are estimates of the marginalized regression functions from Section 2.1. Thus  $\hat{\psi}$  is an adjusted version of the regression-based plug-in estimator  $\int_{\mathcal{T}} w'(t) \hat{m}(t) dt / \int_{\mathcal{T}} w'(t) \hat{\ell}(t) dt$ , where adding inverse-probability-weighted terms to the numerator and denominator is the adjustment required to obtain double robustness.

More standard IV estimators are often computed with a two-stage least squares approach, where in the first stage the treatment variable is regressed on the instrument (and covariates) and then in the second stage the outcome is regressed on the predicted values from the first stage (and covariates). In fact, the weighted average estimator in (7) can also be constructed with a modified version of such a two-stage least squares approach; this may make it more amenable to practical use. Specifically, the following modified two-stage least squares procedure can be used to compute the weighted average estimator (using pseudo-instrument, treatment, and outcome  $w'(Z)$ ,  $A^*$ , and  $Y^*$  respectively):

1. Regress  $A^* = \frac{A - \hat{\lambda}(\mathbf{X}, Z)}{\pi(Z|\mathbf{X})} + \frac{\mathbb{1}\{w'(Z) \neq 0\}}{w'(Z)} \int_{\mathcal{T}} w'(t) \hat{\lambda}(\mathbf{X}, t) dt$  on  $w'(Z)$  without an intercept, and obtain predicted values  $\hat{A}^*$ .
2. Regress  $Y^* = \frac{Y - \hat{\mu}(\mathbf{X}, Z)}{\pi(Z|\mathbf{X})} + \frac{\mathbb{1}\{w'(Z) \neq 0\}}{w'(Z)} \int_{\mathcal{T}} w'(t) \hat{\mu}(\mathbf{X}, t) dt$  on  $\hat{A}^*$ , without an intercept.

Then the coefficient in front of  $\hat{A}^*$  in the second stage equals  $\hat{\psi}$  from (7).

Closed-form estimators are also available even when effect modification is of interest, as long as we project onto linear models of the form  $\gamma(t, \mathbf{v}; \boldsymbol{\psi}) = \mathbf{h}(t, \mathbf{v})^\top \boldsymbol{\psi}$ , for some known mapping  $\mathbf{h} : \mathcal{T} \times \text{supp}(\mathbf{V}) \rightarrow \mathbb{R}^q$ . Specifically, in such cases the estimator  $\hat{\boldsymbol{\psi}}$  defined as the solution to (6) is given by

$$\begin{aligned} \hat{\boldsymbol{\psi}} = & \mathbb{P}_n \left[ \mathbf{g}_1^*(Z, \mathbf{V}) \left\{ \frac{A - \hat{\lambda}(\mathbf{X}, Z)}{\hat{\pi}(Z | \mathbf{X})} \right\} + \int_{\mathcal{T}} \mathbf{g}_1^*(t, \mathbf{V}) \hat{\lambda}(\mathbf{X}, t) dt \right]^{-1} \\ & \times \mathbb{P}_n \left[ \mathbf{g}_2(Z, \mathbf{V}) \left\{ \frac{Y - \hat{\mu}(\mathbf{X}, Z)}{\hat{\pi}(Z | \mathbf{X})} \right\} + \int_{\mathcal{T}} \mathbf{g}_2(t, \mathbf{V}) \hat{\mu}(\mathbf{X}, t) dt \right] \end{aligned}$$

where  $\mathbf{g}_1^*(z, \mathbf{v}) = \frac{\partial}{\partial t} \{\mathbf{h}(t, \mathbf{v}) w(t, \mathbf{v}) \mathbf{h}(t, \mathbf{v})^\top\} |_{t=z}$ , and  $\mathbf{g}_2(z, \mathbf{v}) = \frac{\partial}{\partial t} \{\mathbf{h}(t, \mathbf{v}) w(t, \mathbf{v})\} |_{t=z}$  is as defined in Theorem 2. Closed-form expressions will typically not be available for estimators in general non-linear models; however, since such estimators are still defined as estimating equation-based Z-estimators, they can be computed with standard software (for example, one could use the `optim` function in R). Variance estimation and confidence interval construction will be discussed in the next section.

### 3.3 Asymptotic Theory

In this section we discuss the asymptotic properties of our proposed estimation approach. In particular we show that under weak conditions our estimator is doubly robust and consistent, and that if the nuisance functions are estimated well enough it is asymptotically normal and efficient. Further, asymptotic normality and efficiency are possible even after flexible machine learning-based covariate adjustment. (Our results equally apply to estimators that only solve the efficient influence function estimating equation asymptotically, up to order  $o_p(1/\sqrt{n})$ , such as targeted minimum loss-based estimators.) An analogous result holds for a sample-splitting version of the proposed estimator, as detailed in Section 6 of the Supplementary Materials.

**Theorem 3.** *Assume that:*

1.  $(\hat{\boldsymbol{\psi}}, \hat{\boldsymbol{\eta}}) \xrightarrow{p} (\boldsymbol{\psi}_0, \bar{\boldsymbol{\eta}})$ , where  $\bar{\boldsymbol{\eta}} = (\bar{\pi}, \bar{\lambda}, \bar{\mu})$  with either  $\bar{\pi} = \pi_0$  or  $(\bar{\lambda}, \bar{\mu}) = (\lambda_0, \mu_0)$ .
2. The sequence of functions  $\hat{\boldsymbol{\varphi}}_n = \boldsymbol{\varphi}(\cdot; \hat{\boldsymbol{\psi}}, \hat{\boldsymbol{\eta}})$  and its limit  $\boldsymbol{\varphi}_0 = \boldsymbol{\varphi}(\cdot; \boldsymbol{\psi}_0, \bar{\boldsymbol{\eta}})$  are contained in a Donsker class with  $\|\hat{\boldsymbol{\varphi}}_n - \boldsymbol{\varphi}_0\|_2 = o_p(1)$ .
3. The map  $\boldsymbol{\psi} \rightarrow \mathbb{E}\{\boldsymbol{\varphi}(\mathbf{O}; \boldsymbol{\psi}, \boldsymbol{\eta})\}$  is differentiable at  $\boldsymbol{\psi}_0$  uniformly in  $\boldsymbol{\eta}$  (around  $\bar{\boldsymbol{\eta}}$ ), with invertible derivative matrix  $\mathbf{D}(\boldsymbol{\psi}_0, \boldsymbol{\eta}) \rightarrow \mathbf{D}(\boldsymbol{\psi}_0, \bar{\boldsymbol{\eta}}) \equiv \mathbf{D}_0$ .

*Then the proposed estimator is consistent with rate of convergence*

$$\|\hat{\boldsymbol{\psi}} - \boldsymbol{\psi}_0\| = O_p \left\{ 1/\sqrt{n} + \|\hat{\pi} - \pi_0\|_2 \left( \|\hat{\lambda} - \lambda_0\|_2 + \|\hat{\mu} - \mu_0\|_2 \right) \right\}.$$

*Suppose further that:*

4.  $\|\hat{\pi} - \pi_0\|_2 (\|\hat{\lambda} - \lambda_0\|_2 + \|\hat{\mu} - \mu_0\|_2) = o_p(1/\sqrt{n})$ .

*Then the proposed estimator is asymptotically normal with*

$$\sqrt{n}(\hat{\boldsymbol{\psi}} - \boldsymbol{\psi}_0) \rightsquigarrow N\left(\mathbf{0}, \mathbb{E}[\{\mathbf{D}_0^{-1} \boldsymbol{\varphi}(\mathbf{O}; \boldsymbol{\psi}_0, \boldsymbol{\eta}_0)\}^{\otimes 2}]\right),$$

*and thus semiparametric efficient.*

A proof of Theorem 3 is given in the Supplementary Materials; it follows from standard Z-estimator theory and empirical process results (van der Vaart & Wellner 1996; van der Vaart 2002). The first condition indicates the double robustness of our approach, since some of the nuisance estimators  $\hat{\boldsymbol{\eta}} = (\hat{\pi}, \hat{\lambda}, \hat{\mu})$  can be misspecified. Specifically, as long as either  $\hat{\pi}$  or  $(\hat{\lambda}, \hat{\mu})$  is consistent, then the estimator  $\hat{\boldsymbol{\psi}}$  will be as well. This gives analysts two chances at consistency, and is particularly important in the IV setting since it can be easier to model the instrument density  $\pi$  than the two regression functions  $(\lambda, \mu)$ .

Conditions 2–3 of Theorem 3 are standard regularity conditions for M- and Z-estimators (van der Vaart & Wellner 1996; van der Vaart 2000; van der Vaart 2002). Condition 2 restricts the flexibility of the nuisance estimators (and their limits), but Donsker classes still cover many complex functions. For example, parametric Lipschitz functions are Donsker, but so are many more complicated function types such as infinite-dimensional smooth functions with bounded partial derivatives, VC classes, Sobolev classes, and functions with bounded uniform sectional variation, as well as convex combinations and Lipschitz transformations of any these classes. More discussion and examples can be found in Sections 2.6–2.7 of van der Vaart & Wellner (1996) and Examples 19.6–19.12 of van der Vaart (2000), as well as in Kennedy (2016). Condition 2 is important because it means we do not have to rely on restrictive parametric models to estimate the potentially complicated and high-dimensional nuisance functions  $(\pi, \lambda, \mu)$ , and can instead use more flexible data-adaptive methods. Condition 2 can also be weakened in various ways. For example, the Donsker condition really only needs to hold in a shrinking neighborhood of  $(\boldsymbol{\psi}_0, \bar{\boldsymbol{\eta}})$ , or with high probability as  $n \rightarrow \infty$ ; alternatively we could formulate Condition 2 in terms of weaker entropy or bracketing conditions. Importantly, in Section 6 of the Supplementary Materials,

we discuss how the sample-splitting estimator mentioned earlier (and defined in Section 4 of the Supplementary Materials) can do away with empirical process conditions entirely, allowing arbitrarily flexible nuisance estimators such as random forests. The differentiability in Condition 3 is standard and required to use a delta method-type result (note that the influence function need not be differentiable itself, only its expectation).

Under Conditions 1–3 of Theorem 3, the proposed estimator is consistent with rate of convergence given by  $1/\sqrt{n} + \|\hat{\pi} - \pi_0\|_2(\|\hat{\lambda} - \lambda_0\|_2 + \|\hat{\mu} - \mu_0\|_2)$ . Again the double robustness is apparent since consistency (i.e.,  $\|\hat{\psi} - \psi_0\| = o_p(1)$ ) follows as long as either the instrument density is consistently estimated,  $\|\hat{\pi} - \pi_0\|_2 = o_p(1)$ , or the treatment/outcome regressions are,  $(\|\hat{\lambda} - \lambda_0\|_2 + \|\hat{\mu} - \mu_0\|_2) = o_p(1)$ . Note that this result is agnostic about how well the nuisance functions are estimated, since slow rates on the nuisance estimators  $\hat{\eta}$  will yield slow rates for the parameter of interest  $\hat{\psi}$ . Importantly, this result also shows how double robustness is useful even apart from giving two chances at consistency; in particular, if we estimate the regression functions  $(\lambda, \mu)$  at slow rates, double robustness gives us a chance to obtain faster rates for  $\hat{\psi}$  by consistently estimating  $\pi$ , and vice versa.

For example, if Condition 4 holds so that  $\|\hat{\pi} - \pi_0\|_2(\|\hat{\lambda} - \lambda_0\|_2 + \|\hat{\mu} - \mu_0\|_2) = o_p(1/\sqrt{n})$ , and therefore nuisance estimation is asymptotically negligible, then the estimator  $\hat{\psi}$  is root-n consistent, asymptotically normal, and semiparametric efficient. Note Condition 4 can hold even if the nuisance functions are estimated at slower than parametric root-n rates, so that efficient estimation and valid inference is possible for  $\psi$  even if we use machine learning-based covariate adjustment, via flexible estimation of the nuisance functions  $(\pi, \lambda, \mu)$ . For example, if the nuisance functions  $(\pi, \lambda, \mu)$  are all estimated at faster than  $n^{1/4}$  rates, so that  $\|\hat{\pi} - \pi_0\|_2 = \|\hat{\lambda} - \lambda_0\|_2 = \|\hat{\mu} - \mu_0\|_2 = o_p(n^{-1/4})$ , then Condition 4 holds since  $o_p(n^{-1/4})o_p(n^{-1/4}) = o_p(1/\sqrt{n})$ . Such rates are possible in various flexible models; for instance, under some conditions (Horowitz 2009) generalized additive model estimators can obtain rates of the form  $O_p(n^{-2/5})$ , which is  $o_p(n^{-1/4})$  since  $R_n = O_p(n^{-2/5})$  implies  $n^{1/4}R_n = n^{-3/20}n^{2/5}R_n = O_p(n^{-3/20}) = o_p(1)$ . Condition 4 can also hold if one of  $\pi$  or  $(\lambda, \mu)$  is estimated with a correctly specified parametric model and the other is merely estimated consistently. However, outside of the randomized trial setting, it is typically uncommon to know such detailed parametric structure.

If Condition 4 holds, confidence intervals can be constructed with the bootstrap, or using a direct estimate of the asymptotic variance given in Theorem 3, such as

$$\mathbb{P}_n[\{\hat{\mathbf{D}}^{-1}\varphi(\mathbf{O}; \hat{\psi}, \hat{\eta})\}^{\otimes 2}]$$

where  $\hat{\mathbf{D}} = \mathbb{P}_n\{\partial\varphi(\mathbf{O}; \psi, \hat{\eta})/\partial\psi^T\}|_{\psi=\hat{\psi}}$  is an estimate of the derivative matrix from Condition 3 of Theorem 3. For completeness, we note that if parametric models are used to estimate all three nuisance functions (and either  $\pi$  or  $(\lambda, \mu)$  are correctly modeled) then the bootstrap would still be valid even though Condition 4 fails, since then the contribution from nuisance estimation is asymptotically linear (an analytic expression could also be derived, since  $\hat{\psi}$  and the estimated nuisance parameters solve a large system of estimating equations). However, as we note in the Introduction, parametric nuisance models are likely to be misspecified, except for  $\pi$  if  $Z$  is externally randomized. Inference is somewhat more delicate in a truly doubly robust but nonparametric setting, where one nuisance estimator can be misspecified but  $(\hat{\pi}, \hat{\lambda}, \hat{\mu})$  are estimated flexibly; we leave this to future work.

### 3.4 Model Selection

To this point we have presumed that we have an a priori model  $\gamma(t, \mathbf{v}; \boldsymbol{\psi})$ , which either represents the truth or a low-dimensional projection. When such a priori models are not available, we might instead aim to learn the form of  $\gamma(t, \mathbf{v})$  from data. Thus in this section we propose a doubly robust cross-validation approach for model selection. Model selection is an important issue in causal inference in general, but this is especially the case for the LIV curve, since the latent threshold  $T$  is continuous; thus saturated parametric models are not possible, even when effect modification is not of interest (i.e.,  $\mathbf{V} = \emptyset$ ). In this section we derive the efficient influence function for the risk of a given candidate estimator, and show how it can be used as a doubly robust loss function in the cross-validation framework developed by van der Laan & Dudoit (2003). In particular, our cross-validation approach could be used to select among working models whose complexity increases with sample size, yielding doubly robust yet nonparametric estimators of the local IV curve, in the same spirit as Robins & Rotnitzky (2001) and Kennedy *et al.* (2017).

If we knew the true LIV curve and true distribution  $P$ , we could evaluate the performance of a given estimator  $\hat{\gamma}_k$  by computing the mean squared error risk  $R^*(\hat{\gamma}_k) = \int_{\mathcal{V}} \int_{\mathcal{T}} w(t, \mathbf{v}) \{\gamma(t, \mathbf{v}) - \hat{\gamma}_k(t, \mathbf{v})\}^2 dP(t, \mathbf{v})$ . Alternatively, if we only wanted to compare or rank a set of candidate estimators  $\{\hat{\gamma}_k : k \in \mathcal{K}\}$ , we could use the pseudo-risk

$$R(\hat{\gamma}_k) = \int_{\mathcal{V}} \int_{\mathcal{T}} w(t, \mathbf{v}) \left\{ \hat{\gamma}_k(t, \mathbf{v})^2 - 2\gamma(t, \mathbf{v})\hat{\gamma}_k(t, \mathbf{v}) \right\} dP(t, \mathbf{v}), \quad (8)$$

since  $R(\hat{\gamma}_k) = R^*(\hat{\gamma}_k) - \mathbb{E}\{w(T, \mathbf{V})\gamma(T, \mathbf{V})^2\}$  is simply a shifted version of the mean squared error  $R^*(\hat{\gamma}_k)$ , and the shift does not depend on the candidate estimator  $\hat{\gamma}_k$ . In standard cross-validation it is possible to estimate risk unbiasedly, without worrying about nuisance function estimation; in contrast, in our setting the risk parameter  $R(\hat{\gamma}_k)$  depends on complex nuisance functions via the curve  $\gamma(t, \mathbf{v})$  and the distribution of the threshold  $T$ . Thus estimation of the risk  $R(\hat{\gamma}_k)$  itself requires nuisance estimation, and in fact we can treat  $R(\hat{\gamma}_k)$  as a parameter in its own right, for which we can develop semiparametric theory and estimators. Thus in the next theorem we give the efficient influence function for the risk  $R(\gamma_k)$  for a given fixed candidate  $\gamma_k$ , and go on to show how to use this efficient influence function as a doubly robust loss function for cross-validation-based model selection.

**Theorem 4.** *Consider the same setting and assumptions as in Theorem 2. Under a non-parametric model, the efficient influence function for the risk  $R(\gamma_k)$  defined in (8) for a fixed candidate  $\gamma_k$  is given by  $L(\mathbf{O}; \gamma_k, \boldsymbol{\eta}) - R(\gamma_k)$ , for*

$$\begin{aligned} L(\mathbf{O}; \gamma_k, \boldsymbol{\eta}) = & f_1(Z, \mathbf{V}; \gamma_k) \left\{ \frac{Y - \mathbb{E}(Y \mid \mathbf{X}, Z)}{p(Z \mid \mathbf{X})} \right\} - f_2(Z, \mathbf{V}; \gamma_k) \left\{ \frac{A - \mathbb{E}(A \mid \mathbf{X}, Z)}{p(Z \mid \mathbf{X})} \right\} \\ & + \int_{\mathcal{T}} \left\{ f_1(t, \mathbf{V}; \gamma_k) \mathbb{E}(Y \mid \mathbf{X}, Z = t) - f_2(t, \mathbf{V}; \gamma_k) \mathbb{E}(A \mid \mathbf{X}, Z = t) \right\} dt \end{aligned} \quad (9)$$

where  $\boldsymbol{\eta} = (\pi, \lambda, \mu)$  are the nuisance functions from before, and  $f_1$  and  $f_2$  are defined as

$$f_1(z, \mathbf{v}; \boldsymbol{\psi}) = 2 \frac{\partial}{\partial t} \left\{ w(t, \mathbf{v}) \gamma_k(t, \mathbf{v}) \right\} \Big|_{t=z}, \quad f_2(z, \mathbf{v}; \boldsymbol{\psi}) = \frac{\partial}{\partial t} \left\{ w(t, \mathbf{v}) \gamma_k(t, \mathbf{v})^2 \right\} \Big|_{t=z}.$$

A proof of Theorem 4 is given in the Supplementary Materials, and follows similar logic as the proof of Theorem 2. We also show that  $L(\mathbf{O}; \gamma_k, \boldsymbol{\eta})$  is a doubly robust loss function for the risk  $R(\gamma_k)$  in the sense that  $\mathbb{E}\{L(\mathbf{O}; \gamma_k, \bar{\boldsymbol{\eta}})\} = R(\gamma_k)$  for nuisance function  $\bar{\boldsymbol{\eta}} = (\bar{\pi}, \bar{\lambda}, \bar{\mu})$  as long as either  $\bar{\pi} = \pi_0$  or  $(\bar{\lambda}, \bar{\mu}) = (\lambda_0, \mu_0)$ , and not necessarily both. Thus we can use  $L(\mathbf{O}; \gamma_k, \boldsymbol{\eta})$  as a doubly robust estimating function, similar to how we used  $\varphi(\mathbf{O}; \boldsymbol{\psi}, \boldsymbol{\eta})$  in previous sections. However, since we typically do not have an independent sample to generate candidates  $\hat{\gamma}_k$ , we need to generate them from the same sample in which we estimate risk. Thus we can use sample-splitting to prevent over-fitting.

In particular, we propose using the loss function in (9) for doubly robust model selection following the general approach of van der Laan & Dudoit (2003). This requires some new notation. Let  $\mathbf{S} = (S_1, \dots, S_n)$  denote a random variable independent of the sample that splits the data into training ( $S_i = 0$ ) and test ( $S_i = 1$ ) sets. For example standard  $v$ -fold cross-validation arises by allowing the split variable  $\mathbf{S}$  to take  $v$  different values  $\{\mathbf{S}_1, \dots, \mathbf{S}_v\}$ , each with equal probability  $1/v$ , where  $\sum_i S_{iv} = n/v$  for all  $v$  and  $\sum_v S_{iv} = 1$  for all  $i$ , so that test sets are all of size  $n/v$  and each unit is only used in one test set. Further define  $\mathbb{P}_{\mathbf{s}}^0$  and  $\mathbb{P}_{\mathbf{s}}^1$  as the sub-empirical distributions for the training data  $\{i : S_i = 0\}$  and test data  $\{i : S_i = 1\}$ , respectively, for a given split  $\mathbf{S} = \mathbf{s}$ . Therefore, for example,  $\hat{\boldsymbol{\eta}}(\mathbb{P}_{\mathbf{s}}^0)$  denotes the nuisance function estimates based only on the training set data, and  $\hat{\gamma}_k(\mathbb{P}_{\mathbf{s}}^0)$  denotes the LIV curve estimate based only on the training set data (which also depends on the nuisance function estimates constructed from the training data).

The cross-validation selection approach of van der Laan & Dudoit (2003) is very similar to standard cross-validation, but incorporates extra steps for nuisance function estimation; it proceeds as follows. For a given split  $\mathbf{s}$  and candidate estimator  $\hat{\gamma}_k$ , we first estimate the nuisance functions with the training data to obtain  $\hat{\boldsymbol{\eta}}(\mathbb{P}_{\mathbf{s}}^0)$ , and then estimate the LIV curve with the training data to obtain  $\hat{\gamma}_k(\mathbb{P}_{\mathbf{s}}^0)$ . Then the loss function  $L$  can be evaluated for any observation  $\mathbf{O}_i$  based on these training estimates, and thus we do so on the test data  $\mathbb{P}_{\mathbf{s}}^1$  and compute the average, given by

$$\hat{R}_{\mathbf{s}}(\hat{\gamma}_k) = \int L\left\{\mathbf{o}; \hat{\gamma}_k(\mathbb{P}_{\mathbf{s}}^0), \hat{\boldsymbol{\eta}}(\mathbb{P}_{\mathbf{s}}^0)\right\} d\mathbb{P}_{\mathbf{s}}^1(\mathbf{o}),$$

which we call the estimated risk for candidate  $k$  at the current split  $\mathbf{s}$ . We repeat the above process for each split, average the split-specific risk estimates to get an overall risk estimate for candidate  $k$ , defined as  $\hat{R}(\hat{\gamma}_k) = \mathbb{E}_{\mathbf{s}}\{\hat{R}_{\mathbf{s}}(\hat{\gamma}_k)\}$ , and finally repeat for each candidate  $k \in \mathcal{K}$  and pick the one  $\hat{k}$  that yields the smallest overall risk estimate  $\hat{k} = \arg \min_{k \in \mathcal{K}} \hat{R}(\hat{\gamma}_k)$ . Hence the cross-validation selector can be written as

$$\hat{k} = \arg \min_{k \in \mathcal{K}} \mathbb{E}_{\mathbf{s}} \int L\left\{\mathbf{o}; \hat{\gamma}_k(\mathbb{P}_{\mathbf{s}}^0), \hat{\boldsymbol{\eta}}(\mathbb{P}_{\mathbf{s}}^0)\right\} d\mathbb{P}_{\mathbf{s}}^1(\mathbf{o}). \quad (10)$$

van der Laan & Dudoit (2003) gave conditions under which the risk  $\hat{R}(\hat{\gamma}_{\hat{k}})$  of the above cross-validation selector is asymptotically equivalent to that of an oracle selector given by  $\tilde{k} = \arg \min_{k \in \mathcal{K}} \mathbb{E}_{\mathbf{s}} \int L\left\{\mathbf{o}; \hat{\gamma}_k(\mathbb{P}_{\mathbf{s}}^0), \bar{\boldsymbol{\eta}}\right\} dP(\mathbf{o})$ , along with corresponding finite-sample bounds. One important condition is that the number of candidates does not grow faster than polynomially with sample size, i.e.,  $|\mathcal{K}| \leq n^c$ . We refer to van der Laan & Dudoit (2003) for more details. Note that semiparametric doubly robust post-selection inference is not considered here, but is an important avenue for future work.



## 4 Simulation Study

To explore finite sample properties of our methods, we simulated from a model with

$$\begin{aligned}(Y^0, \mathbf{X}) &\sim N(\mathbf{0}, \mathbf{I}_5), \\ Z \mid \mathbf{X}, Y^0 &\sim \text{TN}\{1.5 \times \text{sign}(\boldsymbol{\alpha}^T \mathbf{X}), 4, (-2, 2)\}, \\ T \mid Z, \mathbf{X}, Y^0 &\sim N(\boldsymbol{\beta}^T \mathbf{X} + Y^0, 1), \\ A &= \mathbb{1}(Z \geq T), \\ Y &= Y^0 + A(\psi T),\end{aligned}$$

where  $\text{TN}\{\mu, \sigma^2, (l, u)\}$  denotes a truncated normal distribution with support  $[l, u]$ ,  $\boldsymbol{\alpha} = (1, 1, -1, -1)^T$ ,  $\boldsymbol{\beta} = (1, -1, -1, 1)^T$ , and  $\psi = 1$ . This setup satisfies the necessary identifying assumptions, and since  $Y^1 - Y^0 = \psi T$  it follows that the LIV curve is linear in the latent threshold  $T$  (regardless of the conditioning set  $\mathbf{V}$ ). It can also be shown that the above setup implies

$$\begin{aligned}\mathbb{E}(A \mid \mathbf{X}, Z) &= \Phi\left(\frac{Z - \boldsymbol{\beta}^T \mathbf{X}}{\sqrt{2}}\right), \\ \mathbb{E}(Y \mid \mathbf{X}, Z) &= \psi \left\{ (\boldsymbol{\beta}^T \mathbf{X}) \Phi\left(\frac{Z - \boldsymbol{\beta}^T \mathbf{X}}{\sqrt{2}}\right) - \sqrt{2} \phi\left(\frac{Z - \boldsymbol{\beta}^T \mathbf{X}}{\sqrt{2}}\right) \right\},\end{aligned}$$

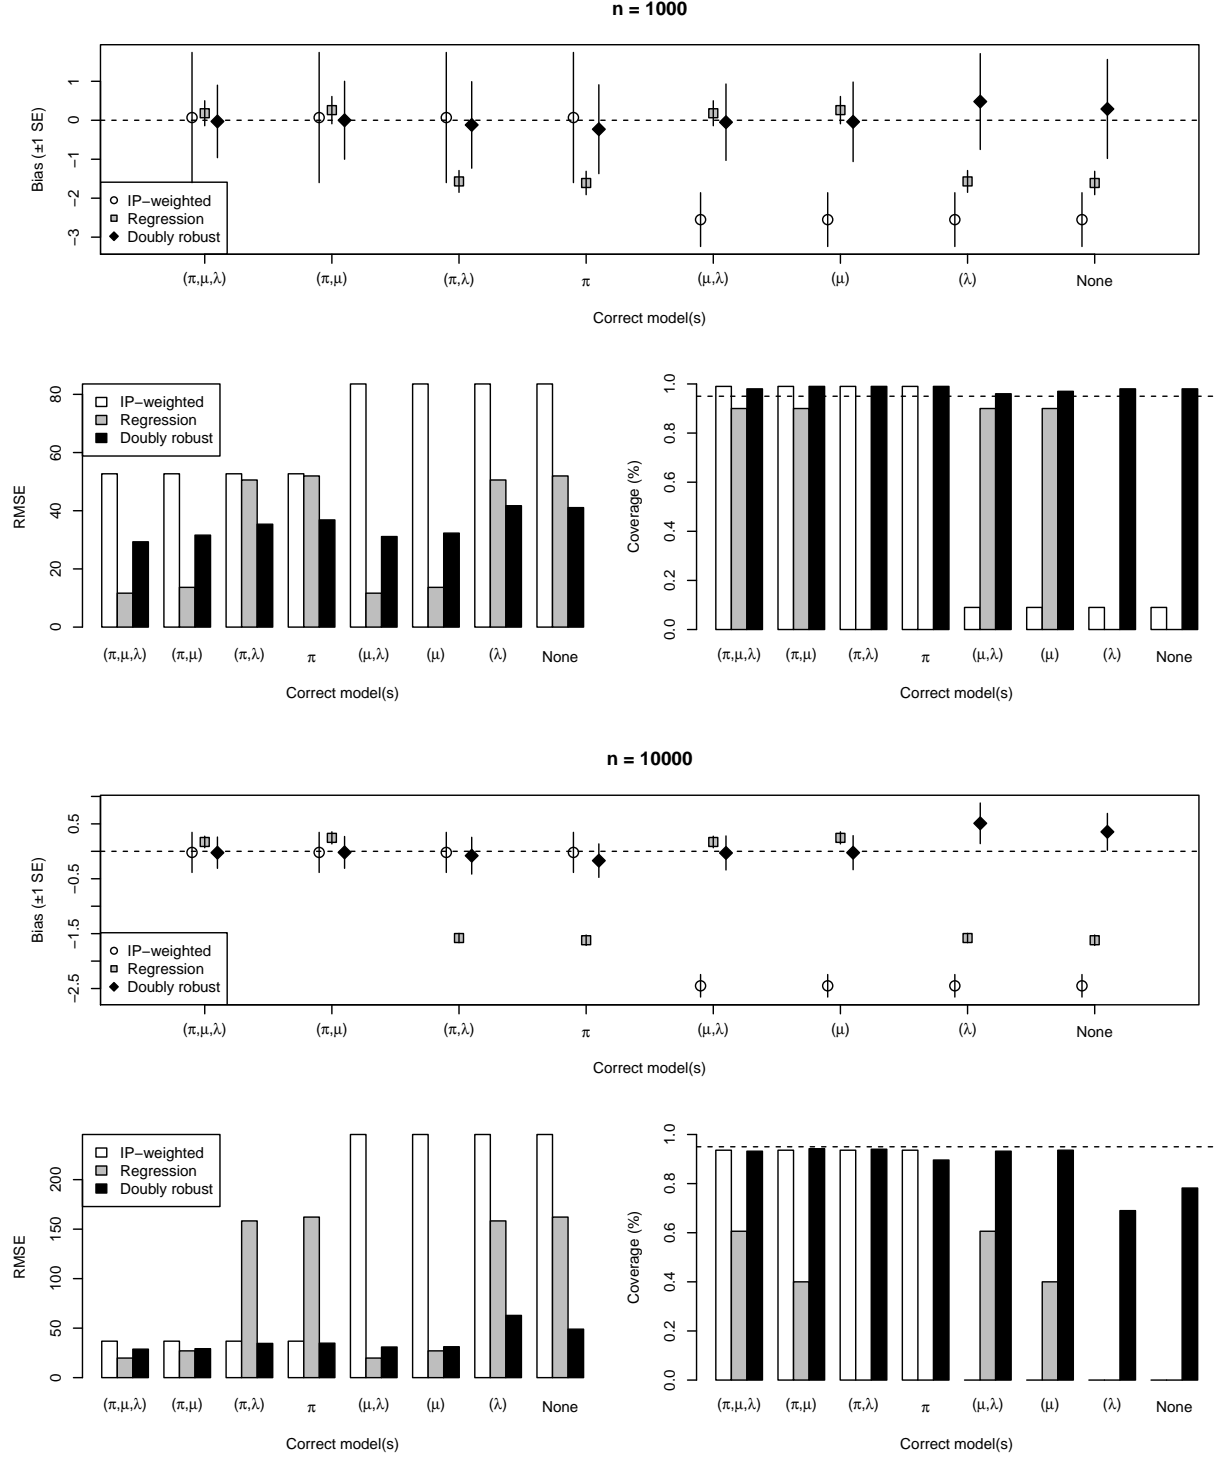
so that the treatment regression  $\lambda$  follows a probit model; the outcome regression  $\mu$  is more complicated, but  $\mathbb{E}(Y \mid \mathbf{X}, Z, A)$  follows a particular generalized additive model.

We fit the proposed estimator, which in this case has a closed form given by

$$\hat{\psi} = \frac{\mathbb{P}_n \left[ \{w(Z) + Zw'(Z)\} \left\{ \frac{Y - \hat{\mu}(\mathbf{X}, Z)}{\hat{\pi}(Z \mid \mathbf{X})} \right\} \right] + \int_{\mathcal{T}} \{w(t) + tw'(t)\} \hat{m}(t) dt}{\mathbb{P}_n \left[ \{2Zw(Z) + Z^2w'(Z)\} \left\{ \frac{A - \hat{\lambda}(\mathbf{X}, Z)}{\hat{\pi}(Z \mid \mathbf{X})} \right\} \right] + \int_{\mathcal{T}} \{2tw(t) + t^2w'(t)\} \hat{\ell}(t) dt}.$$

In particular we considered inverse-probability (IP) -weighted estimators with  $\hat{\mu} = \hat{\lambda} = 0$ , regression estimators with  $\hat{\pi} = \infty$ , and doubly robust estimators that rely on estimates of all of  $(\pi, \lambda, \mu)$ . Note that although we recommend the doubly robust estimators in practice, even these IP and regression estimators are novel and have not been previously proposed, to the best of our knowledge. For the weight function  $w(t)$  we used the density of a  $\text{TN}\{0, 1, (-1.9, 1.9)\}$  truncated normal variable, which roughly matches the marginal distribution of  $Z$ . To misspecify models for  $(\hat{\pi}, \hat{\lambda}, \hat{\mu})$  we used the Kang & Schafer (2007) covariate transformations; to misspecify  $\hat{\lambda}$  we additionally transformed  $Z$  to  $e^Z$  and used a logit rather than probit model. We used maximum likelihood to estimate  $\hat{\pi}$  and  $\hat{\lambda}$ , and a generalized additive model to estimate  $\hat{\mu}$ . Coverage was assessed based on bootstrap confidence intervals, using 100 bootstrap samples. Results (including bias, standard errors, root mean squared error (RMSE), and coverage) are shown in Figure 2.

The weighting estimator gives large bias unless its nuisance estimator  $\hat{\pi}$  is correctly specified, and similarly the regression estimator gives large bias unless  $\hat{\lambda}$  and  $\hat{\mu}$  are correctly specified; however, the doubly robust estimator gives small bias as long as either  $\hat{\pi}$  or



**Figure 2:** Simulation results across 500 simulations: bias (with standard error), RMSE (scaled by  $\sqrt{n}$ ), and bootstrap coverage (based on 100 bootstrap samples).

$(\hat{\lambda}, \hat{\mu})$  are correct. In our setup, even under misspecification, the bias for the doubly robust estimator is smaller than that of the weighting or regression estimators. The weighted estimator is least precise, while the regression estimator is most precise, and can outperform the doubly robust estimator in terms of mean squared error when both  $\hat{\lambda}$  and  $\hat{\mu}$  are estimated well, particularly for smaller sample sizes (indicating a bias-variance trade-off). Coverage roughly coincided with bias: the weighting and regression estimators gave poor coverage unless their nuisance estimators were correctly specified, but the doubly robust estimator gave good coverage as long as either  $\hat{\pi}$  or  $(\hat{\lambda}, \hat{\mu})$  was modeled well. Correctly-specified weighting and doubly robust estimators exhibited some conservative behavior (i.e., greater than 95% coverage) at  $n = 1000$ , and slightly anti-conservative behavior for the larger sample size; but this may be explained by the relatively modest number of bootstrap samples and simulations (which was required to reduce computation time).

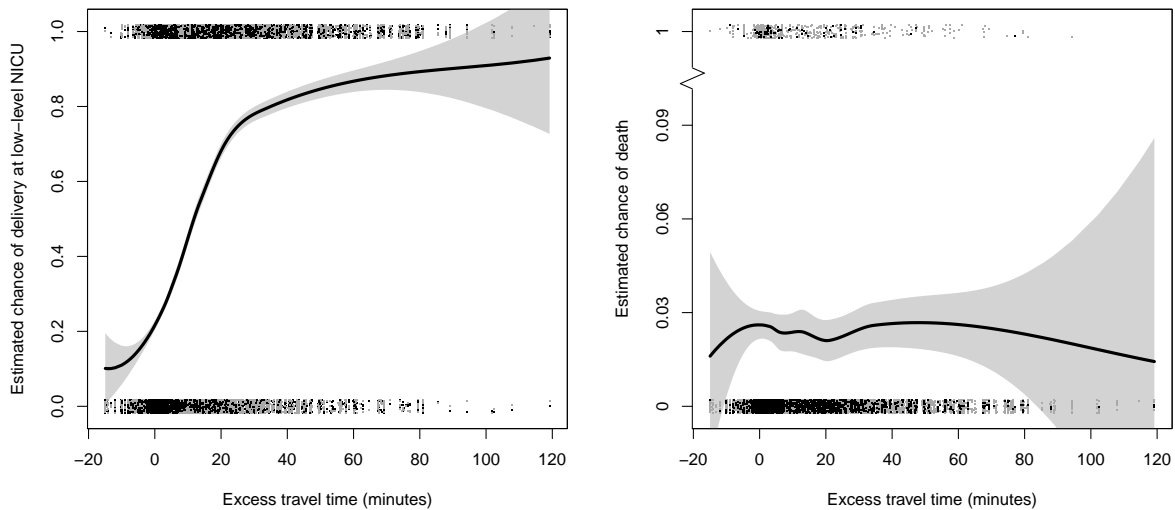
The regression estimator gave slight bias and relatively poor coverage even for large sample sizes; we hypothesize that this is because this estimator relies on a generalized additive model estimator, which may not be smooth enough to guarantee  $\sqrt{n}$  rates and asymptotic linearity. As noted throughout our paper, this is an important motivation for using doubly robust and other influence-function-based estimators, which can attain  $\sqrt{n}$  rates and asymptotic linearity even if their nuisance estimators do not. In the Supplementary Materials we present simulation results based on using flexible methods (generalized additive models and random forests) to estimate all of  $(\pi, \lambda, \mu)$ , not just  $\mu$ .

## 5 Illustration

In this section we apply the proposed methodology to estimate the effects on infant mortality of delivery at hospitals with high- versus low-level neonatal intensive care units (NICUs). Following Lorch *et al.* (2012) and others, we define high-level NICUs as those that are designated as level III by the American Academy of Pediatrics, and that deliver at least 50 low birthweight infants on average per year. Level III units have high technical capacity, providing subspecialist teams, advanced imaging, and the ability for sustained mechanical assisted ventilation. On the other hand, level I-II NICUs are only designed to provide basic care to lower-risk infants. The question of whether and how care at high-level NICUs might impact infant mortality is important from both patient and policy perspectives. For example if high-level units can reduce infant mortality, particularly among high-risk infants, then policies that send high-risk infants to high-level NICUs might be worth pursuing.

To assess potential benefits of delivery at hospitals with high-level units, Lorch *et al.* (2012) collected data on all  $n = 192,078$  premature births in Pennsylvania between 1995 and 2006. Covariate information included data about the infant, such as birthweight and gestational age, as well as about the delivering mother, such as age, race, and measures of socioeconomic status and comorbidities. A full list of covariate information is given in the Supplementary Materials, and more details can be found in Baiocchi *et al.* (2010) and Lorch *et al.* (2012). Importantly, the data are missing some detailed clinical information (e.g., comorbidity severity and lab results) that might explain mothers' deliveries at high- versus low-level hospitals; therefore analyses relying on 'no unmeasured confounding' assumptions

could be suspect. Fortunately, Baiocchi *et al.* (2010) and Lorch *et al.* (2012) identified a potential IV, which is the excess travel time (in minutes) it takes a mother to get to the nearest high- versus low-level NICU. This is a plausible instrument since it affects where mothers deliver (larger values mean mothers have to travel longer to get to high-level units), but it likely does not independently affect infant mortality and is probably not associated with unmeasured confounders that also affect mortality (at least conditional on measured factors like socioeconomic status). More discussion can be found in Baiocchi *et al.* (2010) and Lorch *et al.* (2012). Figure 3 shows loess fits of the unadjusted relationship between instrument and treatment (which is strong), and between instrument and outcome (which is less strong); the raw data points also indicate the marginal distribution of the instrument.



**Figure 3:** Relationship between instrument  $Z$  (excess travel time) and treatment  $A$  (delivery at low-level unit) on the left, and instrument and outcome  $Y$  (infant mortality) on the right. Note that  $y$ -axis scaling differs in the two plots.

We conducted two sets of analyses based on the methodology proposed in previous sections. First we estimated the LIV curve only conditional on the threshold value (so that  $\mathbf{V} = \emptyset$ ), and used the proposed cross-validation approach to select among spline models. Second we estimated how effects vary with birthweight and gestational age, which are two important potential effect modifiers. In both analyses it is first necessary to estimate the nuisance functions, which we did using generalized additive models. To estimate the instrument density  $\pi$ , we used a model previously used by Kennedy *et al.* (2017), in which the density only depends on covariates through the mean and variance functions but is otherwise flexible. Specifically this model assumes  $Z = \pi_1(\mathbf{X}) + \pi_2(\mathbf{X})\epsilon$ , where  $\epsilon$  satisfies  $\mathbb{E}(\epsilon | \mathbf{X}) = 0$  and  $\mathbb{E}(\epsilon^2 | \mathbf{X}) = 1$ , the density  $f_\epsilon$  of  $\epsilon$  is unspecified but smooth, and  $(\pi_1, \pi_2)$  follow generalized additive models with identity and log links, respectively. Thus under this model the conditional density of the instrument is given by  $\pi(z | \mathbf{x}) = f_\epsilon[\{z - \pi_1(\mathbf{x})\}/\pi_2(\mathbf{x})]$ .

In our first analysis we estimated the LIV curve  $\gamma(t)$  using a density-weighted projection based on the marginal density of the instrument, so that  $w(t) = \hat{p}(z = t)$  for  $\hat{p}$  a usual kernel density estimator. A plot of this weight function is given in the Supplementary Materials. The weighted average treatment effects we report should therefore be interpreted as averages of the true LIV curve values, where the averaging weights non-extreme moderate instrument values most highly, according to the weight function plotted in the Supplementary Materials. As noted in Section 3.2, when projecting onto a constant, the resulting weighted average effect also equals the complier average effect in a population where the density of the latent threshold  $T$  equals the weight function  $w^*$ . We used natural cubic splines for  $\gamma_k(t; \psi_k)$  with degrees of freedom  $k \in \{1, 2, 3, 4\}$  selected via cross-validation with two folds, using the proposed doubly robust pseudo-risk estimate  $\hat{R}(\hat{\gamma}_k)$ .

The linear model with  $k = 2$  gave the smallest pseudo-risk ( $-13.7 \times 10^{-6}$ ), much smaller than  $k = 3$  and  $k = 4$ , which led to overfitting ( $-7.3 \times 10^{-6}$  and  $1.5 \times 10^{-3}$ ). The risk under  $k = 2$  was similar to that of the constant effect model with  $k = 1$  ( $-12.6 \times 10^{-6}$ ) and gave very similar estimates. For example, for the linear model the effect estimates range from 9.0 to 8.9 deaths per 1000 births for excess travel times ranging from 0 to 100, and at level 0.05 we cannot reject the hypothesis that the slope parameter equals zero ( $p = 0.98$ ). Table 1 gives estimates and 95% confidence intervals (based on the bootstrap) for three estimators using the constant effect working model  $\gamma(t; \psi) = \psi$ ; the inverse-probability-weighted estimator only relies on estimating the conditional instrument density  $\pi$  (i.e., it plugs in sample averages of  $A$  and  $Y$  for  $\hat{\lambda}$  and  $\hat{\mu}$ ), the regression-based estimator only relies on estimating the treatment and outcome regressions  $(\lambda, \mu)$  (i.e., it plugs in  $\infty$  for  $\hat{\pi}$ ), and the doubly robust estimator is the proposed approach detailed in Section 3.

**Table 1:** Risk difference estimates using weighted average of LIV curve (in terms of deaths prevented per 1000 births).

Method	Est (95% CI)
Inverse-probability-weighted	-4.8 (-17.2, 7.6)
Regression-based	9.2 (6.3, 12.1)
Doubly robust	8.9 (5.4, 12.5)

The proposed doubly robust estimator indicates a mortality benefit (risk difference) of 8.9 fewer deaths per 1000 births due to high-level NICU care (95% CI: 5.4, 12.5), among compliers who could be encouraged by travel time to go to a low-level unit. For comparison, this estimate contrasts with the unadjusted risk difference of -18.6 (-20.0, -17.2), which makes high-level NICUs appear to be harming infants, and a doubly robust no-unmeasured-confounding-based estimate of -0.6 (-2.8, 1.6) for the average treatment effect, which does not give any evidence of benefit. Our estimator yields similar results as the two-stage least squares and matching analyses in Baiocchi *et al.* (2010), but targets a different parameter and relies on weaker assumptions. Our regression and doubly robust estimates were similar, and differed somewhat from weighting, indicating that the conditional density model might be misspecified (although the weighted estimator is also imprecise).

In our second analysis (exploring effect modification by birthweight and gestational age), we projected onto a model in which effects do not vary with the latent threshold (based on the results of our first analysis) but can vary with normal versus low birthweight (2000+ grams versus <2000 grams) and early versus very early gestational age (35–37 weeks versus  $\leq 34$  weeks). Therefore in this analysis we set  $\gamma(t, v; \psi) = \sum_j \psi_j \mathbb{1}(v = j)$  where  $j \in \{1, 2, 3, 4\}$  indexes the four groups. Results are given in Table 2.

**Table 2:** Effect estimates (95% confidence intervals) by birthweight and gestational age.

Birthweight	Gestational age	
	$\leq 34$ wks	35–37 wks
<2000 g	58.5 (52.7, 64.3)	6.1 (2.2, 10.0)
2000+ g	10.1 (2.7, 17.4)	2.4 (-0.8, 5.6)

The largest effect of high-level NICU care was for the highest-risk infants with low birthweight and very early gestational age; in particular, for this group, care at high-level NICUs was estimated to yield 58.5 fewer deaths per 1000 births (95% CI: 52.7, 64.3). Effects in the other two higher-risk groups were relatively similar, with an estimated 6.1 and 10.1 fewer deaths per 1000 births (and both statistically significantly different from zero). For the lower-risk group with higher birthweight and gestational age, care at high-level units was estimated to bring 2.4 fewer deaths per 1000 births (95% CI: -0.8, 5.6), and we cannot reject a null hypothesis of zero average effect. Results were similar but less pronounced when we used less extreme cutoffs for birthweight (2500+ grams versus <2500 grams) and gestational age ( $\leq 35$  weeks versus 36–37 weeks).

## 6 Discussion

In this paper we developed novel semiparametric theory and estimation procedures for a marginal version of the LIV curve, which represents the effect among local compliers who would be encouraged to take treatment at a given threshold value of the instrument but not below. Importantly, in contrast to available methods for estimating the fully conditional LIV curve, our methods have the following advantages: they do not require parametric assumptions (but can still yield parametric root-n rates of convergence), incorporate information about the instrument mechanism, are doubly robust (i.e., still yield consistent estimates under misspecification of either the instrument or treatment/outcome processes), and allow for estimating varying amounts of effect modification. We described the asymptotic properties of our methods under weak empirical process conditions, and also proposed a doubly robust cross-validation approach for model selection. Finally we used the proposed methods to study the effects of care at high-level NICUs on infant mortality, including how such effects are modified by infants’ birthweight and gestational age.

There are a number of direct opportunities for future work based on this research, beyond those areas mentioned in the main text (e.g., extensions to deal with positivity

violations, loss functions beyond  $L_2$ , further development and assessment of the sample-splitting estimator, nonparametric estimation of the local IV curve via kernel-smoothed projections, post-selection inference, etc.). First, it will be very useful to develop computationally efficient software for implementing the proposed methods for general non-linear working models. The methods are computationally demanding due to the need to calculate multiple derivatives and integrals, especially in cases involving complex effect modification. Second, it might be of interest to determine the efficient choice of the weight function  $w(t, \mathbf{v})$  for the case where the working model  $\gamma(t, \mathbf{v}; \boldsymbol{\psi})$  is believed to be the true model. A third area of future work is in the application studying the effects of high-level NICU care, where it would be useful to implement more flexible covariate adjustment (e.g., Super Learner) and more complex models (e.g., exploring threshold effects and continuous effect modification).

## 7 References

1. Abadie, A. Semiparametric instrumental variable estimation of treatment response models. *Journal of Econometrics* **113**, 231–263 (2003).
2. Angrist, J. D., Imbens, G. W. & Rubin, D. B. Identification of causal effects using instrumental variables. *Journal of the American Statistical Association* **91**, 444–455 (1996).
3. Baiocchi, M., Cheng, J. & Small, D. S. Instrumental variable methods for causal inference. *Statistics in Medicine* **33**, 2297–2340 (2014).
4. Baiocchi, M., Small, D. S., Lorch, S. & Rosenbaum, P. R. Building a stronger instrument in an observational study of perinatal care for premature infants. *Journal of the American Statistical Association* **105**, 1285–1296 (2010).
5. Bang, H. & Robins, J. M. Doubly robust estimation in missing data and causal inference models. *Biometrics* **61**, 962–973 (2005).
6. Basu, A., Heckman, J. J., Navarro-Lozano, S. & Urzua, S. Use of instrumental variables in the presence of heterogeneity and self-selection: an application to treatments of breast cancer patients. *Health Economics* **16**, 1133–1157 (2007).
7. Bickel, P. J., Klaassen, C. A., Ritov, Y. & Wellner, J. A. *Efficient and Adaptive Estimation for Semiparametric Models* (Johns Hopkins University Press, 1993).
8. Carneiro, P., Heckman, J. J. & Vytlacil, E. J. Estimating marginal returns to education. *National Bureau of Economic Research Working Paper Series Paper 16474*, 1–32 (2010).
9. Chernozhukov, V. *et al.* Double Machine Learning for Treatment and Causal Parameters. *arXiv preprint arXiv:1608.00060* (2016).
10. Glickman, M. E. & Normand, S.-L. T. The derivation of a latent threshold instrumental variables model. *Statistica Sinica*, 517–544 (2000).
11. Heckman, J. J. Instrumental variables: A study of implicit behavioral assumptions used in making program evaluations. *Journal of Human Resources* **32**, 441–462 (1997).

12. Heckman, J. J. & Vytlacil, E. J. Local instrumental variables and latent variable models for identifying and bounding treatment effects. *Proceedings of the National Academy of Sciences* **96**, 4730–4734 (1999).
13. Heckman, J. J. & Vytlacil, E. J. Structural equations, treatment effects, and econometric policy evaluation. *Econometrica* **73**, 669–738 (2005).
14. Hernán, M. A. & Robins, J. M. Instruments for causal inference: an epidemiologist’s dream? *Epidemiology* **17**, 360–372 (2006).
15. Horowitz, J. L. *Semiparametric and Nonparametric Methods in Econometrics* (Springer, 2009).
16. Imbens, G. W. Instrumental variables: an econometricians perspective. *Statistical Science* **29**, 323–358 (2014).
17. Imbens, G. W. & Angrist, J. D. Identification and estimation of local average treatment effects. *Econometrica* **62**, 467–475 (1994).
18. Kang, J. D. & Schafer, J. L. Demystifying double robustness: A comparison of alternative strategies for estimating a population mean from incomplete data. *Statistical Science* **22**, 523–539 (2007).
19. Kennedy, E. H. Semiparametric theory and empirical processes in causal inference. *In: Statistical Causal Inferences and Their Applications in Public Health Research*, Springer (2016).
20. Kennedy, E. H., Ma, Z., McHugh, M. D. & Small, D. S. Nonparametric methods for doubly robust estimation of continuous treatment effects. *Journal of the Royal Statistical Society: Series B*, (to appear) (2017).
21. Lorch, S. A., Baiocchi, M., Ahlberg, C. E. & Small, D. S. The differential impact of delivery hospital on the outcomes of premature infants. *Pediatrics* **130**, 270–278 (2012).
22. Neugebauer, R. & van der Laan, M. J. Nonparametric causal effects based on marginal structural models. *Journal of Statistical Planning and Inference* **137**, 419–434 (2007).
23. Ogburn, E. L., Rotnitzky, A. & Robins, J. M. Doubly robust estimation of the local average treatment effect curve. *Journal of the Royal Statistical Society: Series B* **77**, 373–396 (2015).
24. Okui, R., Small, D. S., Tan, Z. & Robins, J. M. Doubly robust instrumental variable regression. *Statistica Sinica*, 173–205 (2012).
25. Petersen, M. L., Porter, K. E., Gruber, S., Wang, Y. & van der Laan, M. J. Diagnosing and responding to violations in the positivity assumption. *Statistical Methods in Medical Research* **21**, 31–54 (2010).
26. Robins, J. M. Correcting for non-compliance in randomized trials using structural nested mean models. *Communications in Statistics-Theory and methods* **23**, 2379–2412 (1994).



27. Robins, J. M. The analysis of randomized and non-randomized AIDS treatment trials using a new approach to causal inference in longitudinal studies. *Health Service Research Methodology: A Focus on AIDS*, 113–159 (1989).
28. Robins, J. M. & Rotnitzky, A. Comments on Inference for semiparametric models: Some questions and an answer. *Statistica Sinica* **11**, 920–936 (2001).
29. Rosenblum, M. A. & van der Laan, M. J. Targeted maximum likelihood estimation of the parameter of a marginal structural model. *The International Journal of Biostatistics* **6** (2010).
30. Rubin, D. B. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology* **66**, 688–701 (1974).
31. Stock, J. H., Wright, J. H. & Yogo, M. A survey of weak instruments and weak identification in generalized method of moments. *Journal of Business & Economic Statistics* **20**, 518–529 (2012).
32. Tan, Z. Marginal and nested structural models using instrumental variables. *Journal of the American Statistical Association* **105**, 157–169 (2010).
33. Tan, Z. Regression and weighting methods for causal inference using instrumental variables. *Journal of the American Statistical Association* **101**, 1607–1618 (2006).
34. Tchetgen Tchetgen, E. J. & Vansteelandt, S. Alternative identification and inference for the effect of treatment on the treated with an instrumental variable. *Harvard University Biostatistics Working Paper Series Paper 166* (2013).
35. Tsiatis, A. A. *Semiparametric Theory and Missing Data* (Springer, 2006).
36. van der Laan, M. J. & Dudoit, S. Unified cross-validation methodology for selection among estimators and a general cross-validated adaptive epsilon-net estimator: Finite sample oracle inequalities and examples. *UC Berkeley Division of Biostatistics Working Paper Series Paper 130* (2003).
37. van der Laan, M. J. & Robins, J. M. *Unified Methods for Censored Longitudinal Data and Causality* (Springer, 2003).
38. van der Laan, M. J. & Rubin, D. B. Targeted maximum likelihood learning. *UC Berkeley Division of Biostatistics Working Paper Series Paper 212*, 1–87 (2006).
39. van der Vaart, A. W. *Asymptotic Statistics* (Cambridge University Press, 2000).
40. van der Vaart, A. W. Higher order tangent spaces and influence functions. *Statistical Science* **29**, 679–686 (2014).
41. van der Vaart, A. W. Semiparametric statistics. In: *Lectures on Probability Theory and Statistics*, 331–457 (2002).
42. van der Vaart, A. W. & Wellner, J. A. *Weak Convergence and Empirical Processes* (Springer, 1996).
43. Vytlačil, E. J. Independence, monotonicity, and latent index models: An equivalence result. *Econometrica* **70**, 331–341 (2002).

- 44. Westreich, D. & Cole, S. R. Invited commentary: positivity in practice. *American Journal of Epidemiology* **171**, 674–677 (2010).
- 45. Wooldridge, J. M. *Econometric analysis of cross section and panel data* (MIT Press, 2010).
- 46. Zheng, W. & van der Laan, M. J. Asymptotic theory for cross-validated targeted maximum likelihood estimation. *UC Berkeley Division of Biostatistics Working Paper Series Paper* **273**, 1–58 (2010).

# Supplementary Materials for “Robust causal inference with continuous instruments using the local instrumental variable curve”

## 1 Proof of Theorem 1

First note that

$$Z \perp\!\!\!\perp (A^z, Y^z) \mid \mathbf{X} \implies Z \perp\!\!\!\perp (A^z, Y^{zA^z}) \mid \mathbf{X} \implies Z \perp\!\!\!\perp (A^z, Y^{A^z}) \mid \mathbf{X},$$

where the first implication follows from Assumption 2 (consistency) and the second by Assumption 5 (exclusion restriction).

Therefore

$$\mathbb{E}(Y \mid \mathbf{X}, Z = z) = \mathbb{E}(Y^{A^z} \mid \mathbf{X}, Z = z) = \mathbb{E}(Y^{A^z} \mid \mathbf{X}) = \mathbb{E}\{A^z(Y^1 - Y^0) \mid \mathbf{X}\} + \mathbb{E}(Y^0 \mid \mathbf{X})$$

where the first equality follows from Assumption 2 (consistency), the second since Assumption 4 (unconfoundedness of  $Z$ ) implies  $Z \perp\!\!\!\perp (A^z, Y^{A^z}) \mid \mathbf{X}$  under Assumptions 2 (consistency) and 5 (exclusion restriction) as shown above, and the third by rearranging. By the same logic we have

$$\mathbb{E}(A \mid \mathbf{X}, Z = z) = \mathbb{E}(A^z \mid \mathbf{X}, Z = z) = \mathbb{E}(A^z \mid \mathbf{X}).$$

Assumption 3 (positivity) allows us to write conditional expectations given  $\mathbf{X}$  and  $Z$ .

Therefore, combining the above results gives

$$\mathbb{E}(Y \mid \mathbf{X}, Z = z + \delta) - \mathbb{E}(Y \mid \mathbf{X}, Z = z) = \mathbb{E}\{(A^{z+\delta} - A^z)(Y^1 - Y^0) \mid \mathbf{X}\}$$

and

$$\mathbb{E}(A \mid \mathbf{X}, Z = z + \delta) - \mathbb{E}(A \mid \mathbf{X}, Z = z) = \mathbb{E}(A^{z+\delta} - A^z \mid \mathbf{X}),$$

so that

$$\begin{aligned} \mathbb{E}\{\mathbb{E}(Y \mid \mathbf{X}, Z = z + \delta) - \mathbb{E}(Y \mid \mathbf{X}, Z = z) \mid \mathbf{V}\} &= \mathbb{E}\{(A^{z+\delta} - A^z)(Y^1 - Y^0) \mid \mathbf{V}\} \\ &= \mathbb{E}(Y^1 - Y^0 \mid \mathbf{V}, A^{z+\delta} > A^z)P(A^{z+\delta} > A^z \mid \mathbf{V}) \\ &= \mathbb{E}(Y^1 - Y^0 \mid \mathbf{V}, z < T \leq z + \delta)P(z < T \leq z + \delta \mid \mathbf{V}) \end{aligned}$$

and similarly

$$\begin{aligned}\mathbb{E}\{\mathbb{E}(A \mid \mathbf{X}, Z = z + \delta) - \mathbb{E}(A \mid \mathbf{X}, Z = z) \mid \mathbf{V}\} &= \mathbb{E}(A^{z+\delta} - A^z \mid \mathbf{V}) \\ &= P(A^{z+\delta} > A^z \mid \mathbf{V}) = P(z < T \leq z + \delta \mid \mathbf{V}),\end{aligned}$$

where the first equalities follow by iterated expectation, the second by Assumption 1 (monotonicity), which implies  $A^{z+\delta} - A^z = \mathbb{1}(A^{z+\delta} > A^z)$ , and the third by definition of the latent threshold  $T$ , i.e.,  $\{A^{z+\delta} > A^z\} \iff \{A^{z+\delta} = 1, A^z = 0\} \iff \{T \leq z + \delta, T > z\}$ .

Therefore, letting  $\gamma(\mathbf{v}, t) = \mathbb{E}(Y^1 - Y^0 \mid T = t, \mathbf{V} = \mathbf{v})$ , we have

$$\begin{aligned}\lim_{\delta \rightarrow 0} \frac{1}{\delta} \mathbb{E}\{\mathbb{E}(Y \mid \mathbf{X}, Z = t + \delta) - \mathbb{E}(Y \mid \mathbf{X}, Z = t) \mid \mathbf{V}\} \\ &= \lim_{\delta \rightarrow 0} \frac{1}{\delta} \mathbb{E}(Y^1 - Y^0 \mid \mathbf{V}, t \leq T \leq t + \delta) P(t \leq T \leq t + \delta \mid \mathbf{V}) \\ &= \gamma(t, \mathbf{V}) \lim_{\delta \rightarrow 0} \frac{1}{\delta} \left\{ P(T \leq t + \delta \mid \mathbf{V}) - P(T \leq t \mid \mathbf{V}) \right\} \\ &= \gamma(t, \mathbf{V}) p(T = t \mid \mathbf{V})\end{aligned}$$

and similarly

$$\begin{aligned}\lim_{\delta \rightarrow 0} \frac{1}{\delta} \mathbb{E}\{\mathbb{E}(A \mid \mathbf{X}, Z = t + \delta) - \mathbb{E}(A \mid \mathbf{X}, Z = t) \mid \mathbf{V}\} &= \lim_{\delta \rightarrow 0} \frac{1}{\delta} P(t \leq T \leq t + \delta \mid \mathbf{V}) \\ &= \lim_{\delta \rightarrow 0} \frac{1}{\delta} \left\{ P(T \leq t + \delta \mid \mathbf{V}) - P(T \leq t \mid \mathbf{V}) \right\} \\ &= p(T = t \mid \mathbf{V}),\end{aligned}$$

where the equalities follow by Assumption 7 (continuity). Specifically, the first and third equalities follow by the fact that  $T$  is continuously distributed, with  $p(T = t \mid \mathbf{V}) = \frac{\partial}{\partial t} P(T \leq t \mid \mathbf{V})$ , and the second follows by the continuity of  $\gamma(\mathbf{v}, t)$  in  $t$ .

Therefore

$$\gamma(t, \mathbf{v}) = \frac{\frac{\partial}{\partial z} \mathbb{E}\{\mathbb{E}(Y \mid \mathbf{X}, Z = z) \mid \mathbf{V} = \mathbf{v}\}}{\frac{\partial}{\partial z} \mathbb{E}\{\mathbb{E}(A \mid \mathbf{X}, Z = z) \mid \mathbf{V} = \mathbf{v}\}} \Big|_{z=t}$$

since the denominator is bounded away from zero by Assumption 6 (instrumentation).

## 2 Proof of Theorem 2

In this section we use subscripts to index quantities that depend on the distribution  $P$ ; a zero subscript denotes a quantity evaluated at the true distribution  $P = P_0$ . Thus for example  $\mathbb{E}_P$  denotes expectations under  $P$  and  $\mathbb{E}_0$  denotes expectations under the truth  $P = P_0$ ; similarly  $\psi_P$  denotes the parameter  $\psi = \psi(P)$  as a map  $\psi : P \mapsto \mathbb{R}^q$  and  $\psi_0$  denotes its true value evaluated at  $P_0$ .

In the interest of notational simplicity we make two slight abuses of notation. First, in the main text  $\varphi$  was only proportional to the efficient influence function (i.e.,  $\varphi$  did not include the constant matrix scaling factor, which is unnecessary for solving estimating equations), whereas in this appendix we use  $\varphi$  to denote the full efficient influence function (including the constant matrix scaling factor). Second, the functions  $\mathbf{g}_j(t, \mathbf{v}; \boldsymbol{\psi})$  (for  $j = 1, 2$ ) in the main text are denoted by  $\mathbf{g}'_j(t, \mathbf{v}; \boldsymbol{\psi})$  in this appendix.

We will show that  $\varphi(\mathbf{O}; \boldsymbol{\psi}_P, \boldsymbol{\eta}_P) = \varphi_P(\mathbf{O})$  is the efficient influence function by showing that it is the canonical gradient of the pathwise derivative of  $\boldsymbol{\psi}_P$ , i.e., that  $\varphi_P$  satisfies

$$\left. \frac{\partial \boldsymbol{\psi}_\epsilon}{\partial \epsilon} \right|_{\epsilon=0} = \mathbb{E}_0\{\varphi_0(\mathbf{O}) s_0(\mathbf{O})\}$$

where  $\boldsymbol{\psi}_\epsilon = \boldsymbol{\psi}(P_\epsilon)$  denotes the parameter  $\boldsymbol{\psi}$  evaluated at any regular parametric submodel  $\{P_\epsilon : \epsilon\}$  passing through  $P_0$  at  $\epsilon = 0$ , and  $s_\epsilon(\mathbf{o}_1 | \mathbf{o}_2) = \frac{\partial}{\partial \epsilon^*} \log dP_\epsilon^*(\mathbf{o}_1 | \mathbf{o}_2)|_{\epsilon^*=\epsilon}$  denotes the parametric submodel score for any partition  $(\mathbf{O}_1, \mathbf{O}_2) \subseteq \mathbf{O}$ .

By definition we have

$$\boldsymbol{\psi}_P = \arg \min_{\boldsymbol{\psi} \in \mathbb{R}^q} \int_{\mathcal{V}} \int_{\mathcal{T}} w(t, \mathbf{v}) \left\{ \gamma_P(t, \mathbf{v}) - \gamma(t, \mathbf{v}; \boldsymbol{\psi}) \right\}^2 p(T = t | \mathbf{v}) dt dP(\mathbf{v})$$

and thus

$$\int_{\mathcal{V}} \int_{\mathcal{T}} \left. \frac{\partial \gamma(t, \mathbf{v}; \boldsymbol{\psi})}{\partial \boldsymbol{\psi}} \right|_{\boldsymbol{\psi}=\boldsymbol{\psi}_P} w(t, \mathbf{v}) \left\{ \gamma_P(t, \mathbf{v}) - \gamma(t, \mathbf{v}; \boldsymbol{\psi}_P) \right\} p(T = t | \mathbf{v}) dt dP(\mathbf{v}) = 0.$$

Letting  $m_P(z, \mathbf{v}) = \mathbb{E}_P\{\mathbb{E}_P(Y | \mathbf{X}, Z = z) | \mathbf{V} = \mathbf{v}\}$  and  $m'_P(t, \mathbf{v}) = \frac{\partial}{\partial z} m_P(z, \mathbf{v})|_{z=t}$ , and similarly  $\ell_P(z, \mathbf{v}) = \mathbb{E}_P\{\mathbb{E}_P(A | \mathbf{X}, Z = z) | \mathbf{V} = \mathbf{v}\}$  and  $\ell'_P(t, \mathbf{v}) = \frac{\partial}{\partial z} \ell_P(z, \mathbf{v})|_{z=t}$ , then under the identifying assumptions in the main text we have

$$\gamma_P(t, \mathbf{v}) = \frac{m'_P(t, \mathbf{v})}{\ell'_P(t, \mathbf{v})} \quad \text{and} \quad p(t | \mathbf{v}) = \ell'_P(t, \mathbf{v}).$$

Therefore the restriction above is equivalent to

$$\begin{aligned} 0 &= \int_{\mathcal{V}} \int_{\mathcal{T}} \left. \frac{\partial \gamma(t, \mathbf{v}; \boldsymbol{\psi})}{\partial \boldsymbol{\psi}} \right|_{\boldsymbol{\psi}=\boldsymbol{\psi}_P} w(t, \mathbf{v}) \left\{ m'_P(t, \mathbf{v}) - \gamma(t, \mathbf{v}; \boldsymbol{\psi}_P) \ell'_P(t, \mathbf{v}) \right\} dt dP(\mathbf{v}) \\ &= \int_{\mathcal{V}} \int_{\mathcal{T}} \left\{ \mathbf{g}_2(t, \mathbf{v}; \boldsymbol{\psi}_P) m'_P(t, \mathbf{v}) - \mathbf{g}_1(t, \mathbf{v}; \boldsymbol{\psi}_P) \ell'_P(t, \mathbf{v}) \right\} dt dP(\mathbf{v}) \end{aligned}$$

where  $\mathbf{g}_1$  and  $\mathbf{g}_2$  are  $q$ -vectors (with known functional form not depending on  $P$ ) defined as

$$\mathbf{g}_1(t, \mathbf{v}; \boldsymbol{\psi}) = \mathbf{g}_2(t, \mathbf{v}; \boldsymbol{\psi}) \gamma(t, \mathbf{v}; \boldsymbol{\psi}) \quad \text{and} \quad \mathbf{g}_2(t, \mathbf{v}; \boldsymbol{\psi}) = \left. \frac{\partial \gamma(t, \mathbf{v}; \boldsymbol{\psi}^*)}{\partial \boldsymbol{\psi}^*} \right|_{\boldsymbol{\psi}^*=\boldsymbol{\psi}} w(t, \mathbf{v}).$$

And since the weight satisfies  $w(t, \mathbf{v}) = 0$  for  $t \notin \text{int}(\mathcal{T})$ , integration by parts gives

$$\int_{\mathcal{V}} \int_{\mathcal{T}} \left\{ \mathbf{g}'_1(t, \mathbf{v}; \boldsymbol{\psi}_P) \ell_P(t, \mathbf{v}) - \mathbf{g}'_2(t, \mathbf{v}; \boldsymbol{\psi}_P) m_P(t, \mathbf{v}) \right\} dt dP(\mathbf{v}) = 0,$$

where  $\mathbf{g}'_j(t, \mathbf{v}; \boldsymbol{\psi}) = \partial \mathbf{g}_j(z, \mathbf{v}; \boldsymbol{\psi}) / \partial z|_{z=t}$ .

Evaluating the above at  $P = P_\epsilon$  gives

$$\int_{\mathcal{V}} \int_{\mathcal{T}} \left\{ \mathbf{g}'_1(t, \mathbf{v}; \boldsymbol{\psi}_\epsilon) \ell_\epsilon(t, \mathbf{v}) - \mathbf{g}'_2(t, \mathbf{v}; \boldsymbol{\psi}_\epsilon) m_\epsilon(t, \mathbf{v}) \right\} dt dP_\epsilon(\mathbf{v}) = 0,$$

and differentiating with respect to  $\epsilon$  and evaluating at the truth  $\epsilon = 0$  (using the chain rule) gives

$$\begin{aligned} 0 = & \int_{\mathcal{V}} \int_{\mathcal{T}} \left\{ \frac{\partial \mathbf{g}'_1(t, \mathbf{v}; \boldsymbol{\psi})}{\partial \boldsymbol{\psi}} \Big|_{\boldsymbol{\psi}=\boldsymbol{\psi}_0} \frac{\partial \boldsymbol{\psi}_\epsilon}{\partial \epsilon} \Big|_{\epsilon=0} \ell_0(t, \mathbf{v}) + \mathbf{g}'_1(t, \mathbf{v}; \boldsymbol{\psi}_0) \frac{\partial \ell_\epsilon(t, \mathbf{v})}{\partial \epsilon} \Big|_{\epsilon=0} \right. \\ & \left. - \frac{\partial \mathbf{g}'_2(t, \mathbf{v}; \boldsymbol{\psi})}{\partial \boldsymbol{\psi}} \Big|_{\boldsymbol{\psi}=\boldsymbol{\psi}_0} \frac{\partial \boldsymbol{\psi}_\epsilon}{\partial \epsilon} \Big|_{\epsilon=0} m_0(t, \mathbf{v}) - \mathbf{g}'_2(t, \mathbf{v}; \boldsymbol{\psi}_0) \frac{\partial m_\epsilon(t, \mathbf{v})}{\partial \epsilon} \Big|_{\epsilon=0} \right\} dt dP_0(\mathbf{v}) \\ & + \int_{\mathcal{V}} \int_{\mathcal{T}} \left\{ \mathbf{g}'_1(t, \mathbf{v}; \boldsymbol{\psi}_0) \ell_0(t, \mathbf{v}) - \mathbf{g}'_2(t, \mathbf{v}; \boldsymbol{\psi}_0) m_0(t, \mathbf{v}) \right\} s_0(\mathbf{v}) dt dP_0(\mathbf{v}). \end{aligned}$$

Rearranging, this implies that

$$\begin{aligned} \frac{\partial \boldsymbol{\psi}_\epsilon}{\partial \epsilon} \Big|_{\epsilon=0} = & \mathbf{C}_0^{-1} \int_{\mathcal{V}} \int_{\mathcal{T}} \left[ \mathbf{g}'_1(t, \mathbf{v}; \boldsymbol{\psi}_0) \left\{ \frac{\partial \ell_\epsilon(t, \mathbf{v})}{\partial \epsilon} \Big|_{\epsilon=0} + \ell_0(t, \mathbf{v}) s_0(\mathbf{v}) \right\} \right. \\ & \left. - \mathbf{g}'_2(t, \mathbf{v}; \boldsymbol{\psi}_0) \left\{ \frac{\partial m_\epsilon(t, \mathbf{v})}{\partial \epsilon} \Big|_{\epsilon=0} + m_0(t, \mathbf{v}) s_0(\mathbf{v}) \right\} \right] dt dP_0(\mathbf{v}) \end{aligned}$$

with

$$\mathbf{C}_P = - \int_{\mathcal{V}} \int_{\mathcal{T}} \left\{ \frac{\partial \mathbf{g}'_1(t, \mathbf{v}; \boldsymbol{\psi})}{\partial \boldsymbol{\psi}} \Big|_{\boldsymbol{\psi}=\boldsymbol{\psi}_P} \ell_P(t, \mathbf{v}) - \frac{\partial \mathbf{g}'_2(t, \mathbf{v}; \boldsymbol{\psi})}{\partial \boldsymbol{\psi}} \Big|_{\boldsymbol{\psi}=\boldsymbol{\psi}_P} m_P(t, \mathbf{v}) \right\} dt dP(\mathbf{v}),$$

and

$$\begin{aligned} \frac{\partial \ell_\epsilon(z, \mathbf{v})}{\partial \epsilon} \Big|_{\epsilon=0} &= \frac{\partial}{\partial \epsilon} \mathbb{E}_\epsilon \{ \mathbb{E}_\epsilon(A \mid \mathbf{X}, Z = z) \mid \mathbf{V} = \mathbf{v} \} \Big|_{\epsilon=0} = \frac{\partial}{\partial \epsilon} \int_{\mathcal{W}} \sum_{a \in \{0,1\}} a p_\epsilon(a \mid \mathbf{x}, z) dP_\epsilon(\mathbf{w} \mid \mathbf{v}) \Big|_{\epsilon=0} \\ &= \int_{\mathcal{W}} \sum_{a \in \{0,1\}} a \left\{ s_0(a \mid \mathbf{x}, z) + s_0(\mathbf{w} \mid \mathbf{v}) \right\} p_0(a \mid \mathbf{x}, z) dP_0(\mathbf{w} \mid \mathbf{v}) \\ &= \mathbb{E}_0 \left( \mathbb{E}_0 \left[ A \{ s_0(A \mid \mathbf{X}, Z) + s_0(\mathbf{W} \mid \mathbf{V}) \} \mid \mathbf{X}, Z = z \right] \mid \mathbf{V} = \mathbf{v} \right), \end{aligned}$$

and by the same logic

$$\frac{\partial m_\epsilon(z, \mathbf{v})}{\partial \epsilon} \Big|_{\epsilon=0} = \mathbb{E}_0 \left( \mathbb{E}_0 \left[ Y \{ s_0(Y \mid \mathbf{X}, Z) + s_0(\mathbf{W} \mid \mathbf{V}) \} \mid \mathbf{X}, Z = z \right] \mid \mathbf{V} = \mathbf{v} \right).$$

Now we turn to  $\mathbb{E}_0\{\varphi_0(\mathbf{O})s_0(\mathbf{O})\}$ . The putative efficient influence function  $\varphi_P$  from the main text is given by

$$\begin{aligned}\varphi_P(\mathbf{O}) = \mathbf{C}_P^{-1} & \left[ \mathbf{g}'_1(Z, \mathbf{V}; \psi_P) \left\{ \frac{A - \mathbb{E}_P(A | \mathbf{X}, Z)}{p(Z | \mathbf{X})} \right\} - \mathbf{g}'_2(Z, \mathbf{V}; \psi_P) \left\{ \frac{Y - \mathbb{E}_P(Y | \mathbf{X}, Z)}{p(Z | \mathbf{X})} \right\} \right. \\ & \left. + \int_{\mathcal{T}} \left\{ \mathbf{g}'_1(t, \mathbf{V}; \psi_P) \mathbb{E}_P(A | \mathbf{X}, Z = t) - \mathbf{g}'_2(t, \mathbf{V}; \psi_P) \mathbb{E}_P(Y | \mathbf{X}, Z = t) \right\} dt \right]\end{aligned}$$

(note the inclusion of the scaling matrix  $\mathbf{C}_P^{-1}$ ), and  $s_0(\mathbf{O})$  is the parametric submodel score, which can be decomposed as

$$s_0(\mathbf{O}) = s_0(Y, A | \mathbf{X}, Z) + s_0(Z | \mathbf{X}) + s_0(\mathbf{W} | \mathbf{V}) + s_0(\mathbf{V}).$$

Therefore

$$\begin{aligned}\mathbf{C}_0 \mathbb{E}_0 & \left[ \varphi_0(\mathbf{O}) \{s_0(Y, A | \mathbf{X}, Z) + s_0(Z | \mathbf{X}) + s_0(\mathbf{W} | \mathbf{V}) + s_0(\mathbf{V})\} \right] \\ & = \mathbb{E}_0 \left[ \mathbf{g}'_1(Z, \mathbf{V}; \psi_0) \left\{ \frac{A s_0(A | \mathbf{X}, Z)}{p_0(Z | \mathbf{X})} \right\} - \mathbf{g}'_2(Z, \mathbf{V}; \psi_0) \left\{ \frac{Y s_0(Y | \mathbf{X}, Z)}{p_0(Z | \mathbf{X})} \right\} \right. \\ & \quad \left. + \int_{\mathcal{T}} \left\{ \mathbf{g}'_1(t, \mathbf{V}; \psi_0) \mathbb{E}_0(A | \mathbf{X}, Z = t) - \mathbf{g}'_2(t, \mathbf{V}; \psi_0) \mathbb{E}_0(Y | \mathbf{X}, Z = t) \right\} dt \right. \\ & \quad \left. \times \left\{ s_0(\mathbf{W} | \mathbf{V}) + s_0(\mathbf{V}) \right\} \right] \\ & = \int_{\mathcal{V}} \int_{\mathcal{T}} \left( \mathbf{g}'_1(t, \mathbf{v}; \psi_0) \mathbb{E}_0 \left( \mathbb{E}_0 \left[ A \{s_0(A | \mathbf{X}, Z) + s_0(\mathbf{W} | \mathbf{V})\} \mid \mathbf{X}, Z = t \right] \mid \mathbf{V} = \mathbf{v} \right) \right. \\ & \quad \left. - \mathbf{g}'_2(t, \mathbf{v}; \psi_0) \mathbb{E}_0 \left( \mathbb{E}_0 \left[ Y \{s_0(Y | \mathbf{X}, Z) + s_0(\mathbf{W} | \mathbf{V})\} \mid \mathbf{X}, Z = t \right] \mid \mathbf{V} = \mathbf{v} \right) \right. \\ & \quad \left. + \left\{ \mathbf{g}'_1(t, \mathbf{v}; \psi_0) \mathbb{E}_0(A | \mathbf{X}, Z = t) - \mathbf{g}'_2(t, \mathbf{v}; \psi_0) \mathbb{E}_0(Y | \mathbf{X}, Z = t) \right\} s_0(\mathbf{v}) \right) dt dP_0(\mathbf{v}) \\ & = \int_{\mathcal{V}} \int_{\mathcal{T}} \left[ \mathbf{g}'_1(t, \mathbf{v}; \psi_0) \left\{ \frac{\partial \ell_{\epsilon}(z, \mathbf{v})}{\partial \epsilon} \Big|_{\epsilon=0} + \ell_0(t, \mathbf{v}) s_0(\mathbf{v}) \right\} \right. \\ & \quad \left. - \mathbf{g}'_2(t, \mathbf{v}; \psi_0) \left\{ \frac{\partial m_{\epsilon}(t, \mathbf{v})}{\partial \epsilon} \Big|_{\epsilon=0} + m_0(t, \mathbf{v}) s_0(\mathbf{v}) \right\} \right] dt dP_0(\mathbf{v}) = \mathbf{C}_0 \frac{\partial \psi_{\epsilon}}{\partial \epsilon} \Big|_{\epsilon=0}\end{aligned}$$

where the first equality follows by iterated expectation and the fact that  $\mathbb{E}_0\{s_0(\mathbf{O}_1 | \mathbf{O}_2) | \mathbf{O}_2\} = 0$  for any  $(\mathbf{O}_1, \mathbf{O}_2) \subseteq \mathbf{O}$ , the second follows by iterated expectation, the third follows by iterated expectation and by definition of  $\ell_0$  and  $m_0$  (along with the earlier results for their derivatives with respect to  $\epsilon$ ), and the fourth follows by the expression derived earlier for  $\partial \psi_{\epsilon} / \partial \epsilon |_{\epsilon=0}$ .

Therefore, as long as  $\mathbf{C}_0$  is invertible, we have  $\partial \psi_{\epsilon} / \partial \epsilon |_{\epsilon=0} = \mathbb{E}_0\{\varphi_0(\mathbf{O})s_0(\mathbf{O})\}$  and thus  $\varphi_P(\mathbf{O})$  is the efficient influence function.

### 3 Double robustness of efficient influence function $\varphi$

Here we will show that  $\mathbb{E}\{\varphi(\mathbf{O}; \boldsymbol{\psi}, \bar{\pi}, \bar{\lambda}, \bar{\mu})\} = 0$  as long as either

$$\bar{\pi} = \pi_0 \quad \text{or} \quad (\bar{\lambda}, \bar{\mu}) = (\lambda_0, \mu_0).$$

In this section expectations  $\mathbb{E} = \mathbb{E}_0$  and parameters  $\boldsymbol{\psi} = \boldsymbol{\psi}_0$  are evaluated under  $P_0$ , but we drop the subscript for notational convenience.

First note that

$$\begin{aligned} \mathbf{C}_0 \mathbb{E}\{\varphi(\mathbf{O}; \boldsymbol{\psi}, \bar{\pi}, \bar{\lambda}, \bar{\mu})\} &= \left[ \mathbf{g}'_1(Z, \mathbf{V}; \boldsymbol{\psi}) \left\{ \frac{A - \bar{\lambda}(\mathbf{X}, Z)}{\bar{\pi}(Z | \mathbf{X})} \right\} - \mathbf{g}'_2(Z, \mathbf{V}; \boldsymbol{\psi}) \left\{ \frac{Y - \bar{\mu}(\mathbf{X}, Z)}{\bar{\pi}(Z | \mathbf{X})} \right\} \right. \\ &\quad \left. + \int_{\mathcal{T}} \left\{ \mathbf{g}'_1(t, \mathbf{V}; \boldsymbol{\psi}) \bar{\lambda}(\mathbf{X}, t) - \mathbf{g}'_2(t, \mathbf{V}; \boldsymbol{\psi}) \bar{\mu}(\mathbf{X}, t) \right\} dt \right] \\ &= \mathbb{E} \left[ \mathbf{g}'_1(Z, \mathbf{V}; \boldsymbol{\psi}) \left\{ \frac{\lambda_0(\mathbf{X}, Z) - \bar{\lambda}(\mathbf{X}, Z)}{\bar{\pi}(Z | \mathbf{X})} \right\} - \mathbf{g}'_2(Z, \mathbf{V}; \boldsymbol{\psi}) \left\{ \frac{\mu_0(\mathbf{X}, Z) - \bar{\mu}(\mathbf{X}, Z)}{\bar{\pi}(Z | \mathbf{X})} \right\} \right. \\ &\quad \left. + \int_{\mathcal{T}} \left\{ \mathbf{g}'_1(t, \mathbf{V}; \boldsymbol{\psi}) \bar{\lambda}(\mathbf{X}, t) - \mathbf{g}'_2(t, \mathbf{V}; \boldsymbol{\psi}) \bar{\mu}(\mathbf{X}, t) \right\} dt \right] \\ &= \mathbb{E} \int_{\mathcal{Z}} \left[ \mathbf{g}'_1(t, \mathbf{V}; \boldsymbol{\psi}) \left\{ \lambda_0(\mathbf{X}, t) - \bar{\lambda}(\mathbf{X}, t) \right\} - \mathbf{g}'_2(t, \mathbf{V}; \boldsymbol{\psi}) \left\{ \mu_0(\mathbf{X}, t) - \bar{\mu}(\mathbf{X}, t) \right\} \right] \frac{\pi_0(t | \mathbf{X})}{\bar{\pi}(t | \mathbf{X})} dt \\ &\quad + \int_{\mathcal{T}} \left\{ \mathbf{g}'_1(t, \mathbf{V}; \boldsymbol{\psi}) \bar{\lambda}(\mathbf{X}, t) - \mathbf{g}'_2(t, \mathbf{V}; \boldsymbol{\psi}) \bar{\mu}(\mathbf{X}, t) \right\} dt \\ &= \mathbb{E} \int_{\mathcal{T}} \left[ \mathbf{g}'_1(t, \mathbf{V}; \boldsymbol{\psi}) \left\{ \lambda_0(\mathbf{X}, t) - \bar{\lambda}(\mathbf{X}, t) \right\} - \mathbf{g}'_2(t, \mathbf{V}; \boldsymbol{\psi}) \left\{ \mu_0(\mathbf{X}, t) - \bar{\mu}(\mathbf{X}, t) \right\} \right] \left\{ \frac{\pi_0(t | \mathbf{X})}{\bar{\pi}(t | \mathbf{X})} - 1 \right\} dt \\ &\quad + \int_{\mathcal{T}} \left\{ \mathbf{g}'_1(t, \mathbf{V}; \boldsymbol{\psi}) \lambda_0(\mathbf{X}, t) - \mathbf{g}'_2(t, \mathbf{V}; \boldsymbol{\psi}) \mu_0(\mathbf{X}, t) \right\} dt \end{aligned}$$

where the first equality is true by definition, the second and third holds by iterated expectation given  $(\mathbf{X}, Z)$  and  $\mathbf{X}$ , respectively, and the last follows after rearranging and since  $\mathbf{g}'_1 = \mathbf{g}'_2 = 0$  for  $t \notin \text{int}(\mathcal{T})$ .

Therefore if  $\bar{\pi} = \pi_0$  or  $(\bar{\lambda}, \bar{\mu}) = (\lambda_0, \mu_0)$  then  $\mathbf{C}_0 \mathbb{E}\{\varphi(\mathbf{O}; \boldsymbol{\psi}, \bar{\pi}, \bar{\lambda}, \bar{\mu})\}$  equals

$$\begin{aligned} &\int_{\mathcal{T}} \mathbb{E} \left\{ \mathbf{g}'_1(t, \mathbf{V}; \boldsymbol{\psi}) \lambda_0(\mathbf{X}, t) - \mathbf{g}'_2(t, \mathbf{V}; \boldsymbol{\psi}) \mu_0(\mathbf{X}, t) \right\} dt \\ &= \int_{\mathcal{V}} \int_{\mathcal{T}} \mathbf{g}'_1(t, \mathbf{v}; \boldsymbol{\psi}) \ell_0(t, \mathbf{v}) - \mathbf{g}'_2(t, \mathbf{v}; \boldsymbol{\psi}) m_0(t, \mathbf{v}) \Big\} dt dP(\mathbf{v}) = 0 \end{aligned}$$

where the first equality follows by iterated expectation given  $\mathbf{V}$  (and by the definitions of  $\ell$  and  $m$  from the previous section), and the second follows from the restriction given in the previous section (after using integration by parts).



## 4 Sample-splitting estimator

For simplicity, the estimator presented in the main text solves an estimating equation that depends on nuisance estimates  $\hat{\boldsymbol{\eta}} = (\hat{\pi}, \hat{\lambda}, \hat{\mu})$  constructed from the entire sample. The asymptotic normality of such estimators requires empirical process (e.g., Donsker) conditions that restrict the complexity of the nuisance functions and their estimators. Unfortunately, these conditions may not be satisfied for very adaptive machine learning methods, such as random forests; however, such conditions are not necessary for estimators that use sample splitting to separate nuisance function estimation from estimating equation evaluation. We present such an estimator here.

We rely on the notation presented in Section 3.4 of the main text. Again let  $\mathbf{S} = (S_1, \dots, S_n)$  denote a random variable independent of the sample that splits the data into training ( $S_i = 0$ ) and test ( $S_i = 1$ ) sets. Here we suppose the split variable  $\mathbf{S}$  takes  $v$  different values  $\{\mathbf{S}_1, \dots, \mathbf{S}_v\}$ , each with equal probability  $1/v$ , where  $\sum_i S_{iv} = n/v$  for all  $v$  and  $\sum_v S_{iv} = 1$  for all  $i$ , so that test sets are all of size  $n/v$  and each unit is only used in one test set. Define  $\mathbb{P}_{\mathbf{S}}^0$  and  $\mathbb{P}_{\mathbf{S}}^1$  as the sub-empirical distributions for the training data  $\{i : S_i = 0\}$  and test data  $\{i : S_i = 1\}$ , respectively, for a given split  $\mathbf{S} = \mathbf{s}$ . Therefore, for example,  $\hat{\boldsymbol{\eta}}(\mathbb{P}_{\mathbf{S}}^0)$  denotes the nuisance function estimates based only on the training set data. Then the proposed sample-splitting version of our estimator is given by the solution to

$$\mathbb{E}_{\mathbf{S}} \mathbb{P}_{\mathbf{S}}^1[\boldsymbol{\varphi}\{\mathbf{O}; \boldsymbol{\psi}, \hat{\boldsymbol{\eta}}(\mathbb{P}_{\mathbf{S}}^0)\}] = \mathbb{E}_{\mathbf{S}} \int \boldsymbol{\varphi}\{\mathbf{o}; \boldsymbol{\psi}, \hat{\boldsymbol{\eta}}(\mathbb{P}_{\mathbf{S}}^0)\} d\mathbb{P}_{\mathbf{S}}^1(\mathbf{o}) = \mathbf{0},$$

with  $\boldsymbol{\varphi}$  as defined in the main text.

## 5 Proof of Theorem 3

Since  $\hat{\boldsymbol{\psi}}$  is a Z-estimator, Theorem 3 follows directly from Theorem 5.31 of van der Vaart (2000), together with the fact that  $\mathbf{C}_0 \mathbb{P}\{\boldsymbol{\varphi}(\mathbf{O}; \boldsymbol{\psi}, \hat{\boldsymbol{\eta}})\} = \mathbf{C}_0 \int \boldsymbol{\varphi}(\mathbf{o}; \boldsymbol{\psi}, \hat{\boldsymbol{\eta}}) dP(\mathbf{o})$  equals

$$\begin{aligned} \mathbb{P} \int_{\mathcal{T}} & \left[ \mathbf{g}'_1(t, \mathbf{V}; \boldsymbol{\psi}) \left\{ \lambda_0(\mathbf{X}, t) - \hat{\lambda}(\mathbf{X}, t) \right\} - \mathbf{g}'_2(t, \mathbf{V}; \boldsymbol{\psi}) \left\{ \mu_0(\mathbf{X}, t) - \hat{\mu}(\mathbf{X}, t) \right\} \right] \left\{ \frac{\pi_0(t | \mathbf{X})}{\hat{\pi}(t | \mathbf{X})} - 1 \right\} dt \\ &= \mathbb{P} \left[ \left\{ \mathbf{g}'_1(\lambda_0 - \hat{\lambda}) - \mathbf{g}'_2(\mu_0 - \hat{\mu}) \right\} (\pi_0 - \hat{\pi}) / \hat{\pi} \pi_0 \right] \\ &\leq C \| \mathbf{g}'_1(\lambda_0 - \hat{\lambda}) - \mathbf{g}'_2(\mu_0 - \hat{\mu}) \|_2 \| \pi_0 - \hat{\pi} \|_2 \\ &= O_p \left\{ \left( \| \lambda_0 - \hat{\lambda} \|_2 + \| \mu_0 - \hat{\mu} \|_2 \right) \| \pi_0 - \hat{\pi} \|_2 \right\} \end{aligned}$$

where the first line follows by iterated expectation (first given  $\mathbf{X}$  and  $Z$ , then given  $\mathbf{X}$ ), the second by multiplying and dividing by  $\pi(t | \mathbf{X})$ , the inequality follows by Cauchy-Schwarz (i.e.,  $\mathbb{P}(fg) \leq \|f\|_2 \|g\|_2$ ) and boundedness of  $1/\hat{\pi}\pi_0$  (from the Donsker condition and positivity), and the last line by the triangle inequality and boundedness of  $\mathbf{g}'_1$  and  $\mathbf{g}'_2$ .

## 6 Avoiding Donsker conditions with sample-splitting

For the asymptotic properties of the sample-splitting estimator presented in Section 4 of the Supplementary Materials, we refer to Chernozhukov *et al.* (2016), who give a thorough analysis of the conditions required for asymptotic linearity. Conditions (i)–(v) of their Assumption 2.1 are standard regularity and identifiability conditions in Z-estimator problems, which can be verified for sufficiently smooth working models  $\gamma(t, \mathbf{v}; \boldsymbol{\psi})$  and estimating functions with sufficiently well-separated solutions in expectation. Conditions (i)–(vi) of their Assumption 2.2 will be satisfied under the rate conditions given in our Assumption 4 of Theorem 3 in the main text, along with sufficient smoothness of the working models  $\gamma(t, \mathbf{v}; \boldsymbol{\psi})$ . Notably, only rate conditions are required for the nuisance estimators  $\hat{\boldsymbol{\eta}}$ , and not Donsker or other regularity conditions. Zheng & van der Laan (2010) give similar results for sample-splitting-based targeted maximum likelihood estimators, relying on higher-level conditions.

## 7 Proof of Theorem 4 & double robustness of loss $L$

After replacing  $\mathbf{g}_j$  with  $f_j$ , these proofs follow precisely the same logic as the proofs given in previous sections of Theorem 2 and of double robustness of the efficient influence function  $\boldsymbol{\varphi}$ , respectively, and so are omitted.

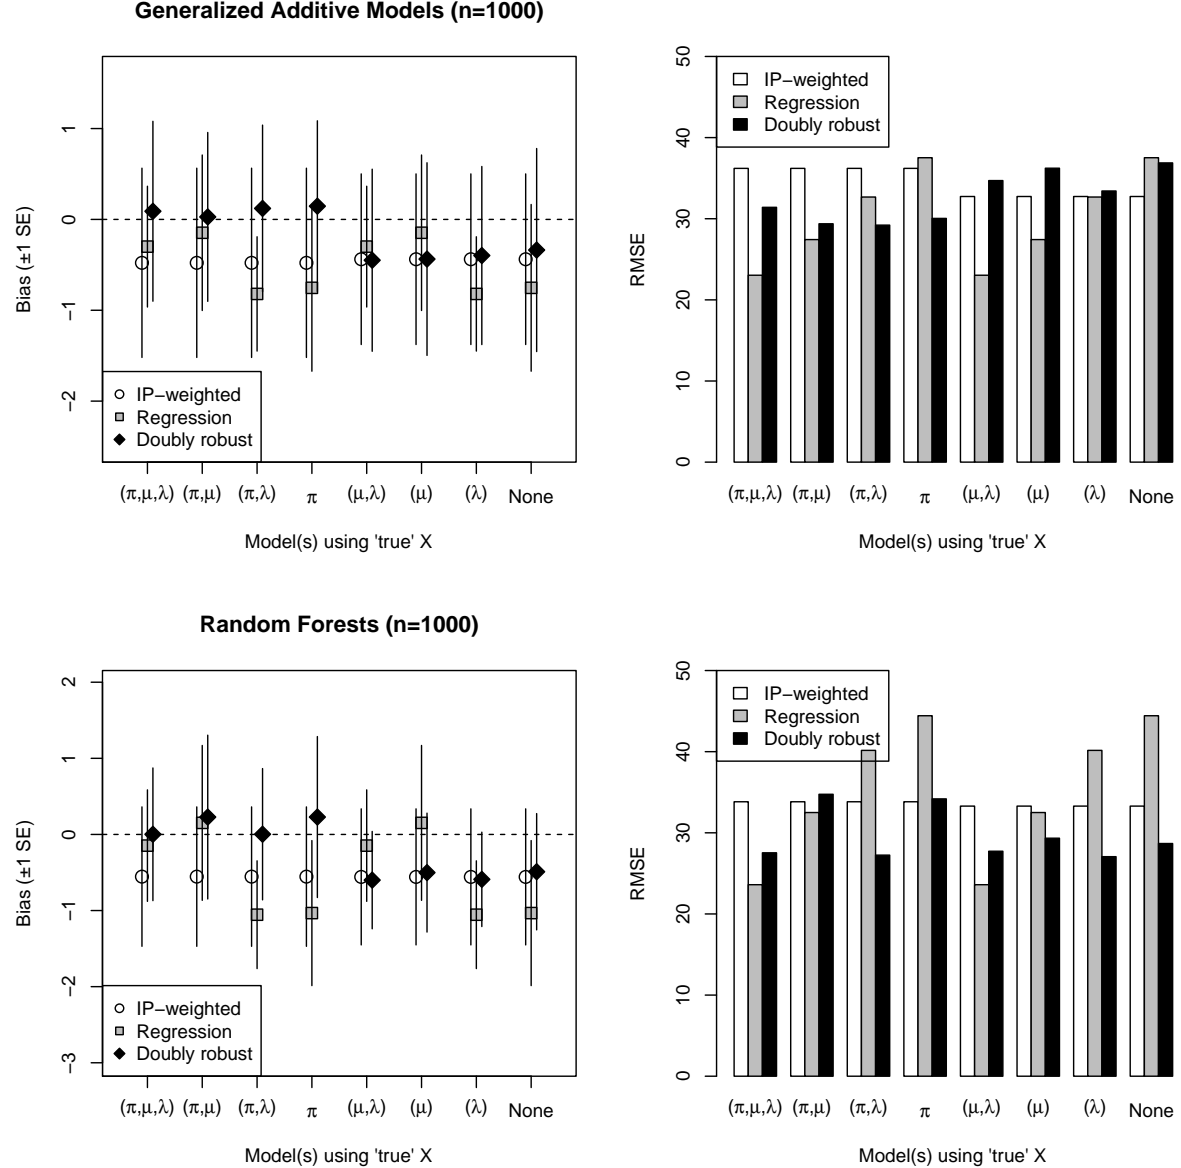
## 8 Additional simulation results

In the simulations in the main text, we used correct models for the nuisance functions (parametric models for  $(\pi, \lambda)$  and a generalized additive model for  $\mu$ ), with misspecification coming from covariate transformations. In this section we present results using generalized additive models and random forests for all nuisance estimators, since it is of interest to explore the performance of our approach when all nuisance functions are estimated without relying on parametric models. The nuisance functions  $(\lambda, \mu)$  are just regression functions so we simply use generalized additive models and random forests as usual there; since  $\pi$  is a density we first use generalized additive models or random forests to estimate the conditional expectation  $\mathbb{E}(Z \mid \mathbf{X})$ , and then use a kernel density estimator to estimate the density depending on the sign of the regression function. Note that these approaches are not necessarily correctly specified (except for estimating  $\mu$  with a generalized additive model, which actually holds). As in the main text, we compare performance using both the “true” covariates  $\mathbf{X}$  and the “misspecified” covariates based on the Kang & Schafer (2007) transformations.

Results are given in Figure 1 below. Since the models were not necessarily correctly specified, the findings are harder to generalize. The inverse-probability-weighted estimator gave some bias in all settings, which suggests that the density estimation procedure could be improved. The regression estimator tended to give small bias when it relied on the true covariates for the  $\mu$  estimator, as might be expected. The doubly robust estimator gave

smallest bias when it relied on the true covariates for the  $\pi$  estimator; it is unclear why there was bias otherwise.

The findings were slightly more consistent for RMSE, especially when using random forests. However, for both generalized additive model and random forests approaches, the doubly robust estimator gave smallest average RMSE across simulation settings, with the inverse-probability-weighted estimator yielding highest average RMSE, and the regression estimator somewhere in between.



**Figure 4:** Simulation results across 500 simulations: bias (with standard error), RMSE (scaled by  $\sqrt{n}$ ), and bootstrap coverage (based on 100 bootstrap samples).

## 9 Details about empirical example

### 9.1 List of adjustment covariates

The data analysis in Section 5 of the main text relied on the following 16 adjustment covariates, as discussed in more detail by Baiocchi et al. (2010) and Lorch et al. (2012).

Information about the zip code in which the mother lives (6):

- median income, percentage below poverty, median home value, percent with high school degree, percent with college degree, percent who rent versus own home.


Information about the mother (8):

- age, diabetes status, month prenatal care was started, number of times previously given birth, whether multiple deliveries, education level (8th grade or less, some high school, high school graduate, some college, college graduate, or more than college), mother's race (White, Black, Asian/Pacific Islander, or other), insurance type (fee for service, HMO, federal/state, other, or uninsured).

Information about the infant (2):

- birthweight, gestational age.

### 9.2 Plot of estimated density of $Z$ (i.e., weight function $w$ )



**Figure 5:** Kernel estimate of the marginal density  $p(z = t)$  of the instrument  $Z$ , also used as the weight function  $w(t) = \hat{p}(z = t)$  in the analysis in Section 5 of the main text.